

Career Concerns in Collective Decision-Making: The Federal Open Market Committee*

Matias Iaryczower
Princeton

Gabriel Lopez Moctezuma
Caltech

Paola Moscardello
Yale[†]

Jan, 2026

Abstract

In this paper, we quantify the distortions induced by career concerns within the Federal Open Market Committee (FOMC). We combine a structural approach with an unanticipated change in the information available to the public about internal committee deliberations. We show that – given the policy preferences of Fed Presidents and Board Governors serving in the FOMC – agents’ incentives to appear competent and unbiased outweigh the distortions induced by anti-pandering and conformity. Relative to a counterfactual with no reputational considerations, career concerns improve the welfare of an unbiased principal. Given our estimates of career concerns, Transparency improves welfare relative to an Opaque regime in which internal deliberations are not made public. In a counterfactual exercise, we show that greater heterogeneity in regional shocks reduces conformity but increases policy errors under Transparency.

*We thank audiences at Arizona, Caltech, Princeton, and the Polecon UK Webinar for helpful comments. We are grateful to Camila Blanes for outstanding research assistance.

[†]Matias Iaryczower: Department of Politics, Princeton University, email: miaryc@princeton.edu; Gabriel Lopez Moctezuma: HSS, California Institute of Technology, email: glopezmo@caltech.edu; Paola Moscardello: Cowles Foundation for Research in Economics, Yale University, email: paola.moscardello@yale.edu.

1 Introduction

Since the pioneering work of Holmström (1982), the theoretical and empirical literature in economics has made substantial progress in understanding how *career concerns* distort decision-makers’ behavior. In a nutshell, agents’ career prospects depend on the reputation they build with outside observers, who evaluate their performance based only on partial information about the underlying fundamentals of the problem faced by the decision-makers. As a result, agents who care enough about *their* future have incentives to manage their reputation, even at the expense of implementing policies they believe to be superior.¹

The study of distortions induced by career concerns has received particular attention in political economy, where contracts between principals and agents are severely limited. This perspective has naturally extended the literature to settings in which policy-making is in the hands of committees, a setup ubiquitous in public institutions (Congress, the FOMC, FDA advisory committees, etc). This effort has generated a wealth of theoretical and empirical advances. To the best of our knowledge, however, we still lack a systematic assessment of the costs of distortions induced by career concerns in collective decision-making bodies.

In this paper, we study this problem in the Federal Open Market Committee (FOMC), the body responsible for setting monetary policy in the United States. To evaluate the welfare effects of career concerns, we adopt a structural estimation approach, using data from internal FOMC deliberations between 1970 and 2008. This period spans a “natural experiment” – first used by Meade and Stasavage (2008) – that provides individual-level data from two distinct regimes: an *Opaque* regime, in which members believed their deliberations would remain secret, and a *Transparent* regime, in which they were aware their remarks would be made public. Exploiting this regime change, together with the equilibrium restrictions implied by the model, we (i) estimate the intrinsic value members attach to reputation, (ii) identify the behavioral distortions induced by career concerns, (iii) quantify the probability of correct recommendations, and (iv) assess the performance of the FOMC under counterfactual committee compositions and economic environments.

The FOMC provides a useful case study for three reasons. First, the FOMC

¹These adverse-selection incentives differ from those in moral-hazard settings, where monitoring by the principal can induce higher levels of effort by the agent.

reigns over a crucial policy domain. Monetary policy is one of the most important instruments to regulate the economy, with broad implications for virtually all economic agents. Understanding which institutional arrangements foster better policy decisions is therefore of paramount importance. Second, there is robust reduced-form evidence of non-negligible reputational considerations in the FOMC (Meade and Stasavage (2008); Swank, Visser, and Swank (2008); Hansen, McMahon, and Prat (2018)), based on the unexpected shift from the Opaque to the Transparent regime. Third, the FOMC features a mixed membership, consisting of seven members of the Board of *Governors* of the Federal Reserve System and twelve regional Reserve Bank *Presidents*. This mixed system allows us to exploit differences in how members with distinct reputational incentives respond to the same economic fundamentals.

From both a theoretical and empirical standpoint, the fact that monetary policy decisions are determined in a committee introduces challenges that are absent in individual decision-making contexts. For a single decision-maker, outsiders' expectations about the optimal action are determined solely by their prior beliefs. In this setting, career concerns may lead the agent either to align excessively with the principal's prior—resulting in *pandering* (Canes-Wrone, Herron, and Shotts (2001); Maskin and Tirole (2004); Prat (2005)) – or to contradict it excessively, resulting in *anti-pandering* (Levy (2004)). When decision-making rests with a committee, outsiders can infer information about the optimal policy from the actions of other members. As a result, the signaling value of an individual's decision depends on the behavior of fellow members. In this context, agents may have incentives either to present a unified front—*conforming* or *herding*—or, conversely, to differentiate themselves from others (Scharfstein and Stein (1990); Zwiebel (1995); Prendergast and Stole (1996); Ottaviani and Sørensen (2001); Visser and Swank (2007); Levy (2007)).

The collaborative nature of decision-making in the FOMC poses significant challenges for identifying the effects of career concerns on outcomes and welfare. In this setting, members' reputations depend not only on their own actions but also on the behavior of others. Moreover, members respond to information that is unobserved by the researcher and correlated across members, making it difficult to separately identify the roles of private information, individual preferences, and career concerns using reduced-form methods. Our structural model—together with the unanticipated regime shift—addresses these challenges by disentangling the effects of preferences, information, and career concerns on members' behavior. This framework also allows

us to conduct counterfactual analyses to analyze how behavior in the FOMC would differ under the Opaque and Transparent regimes, and in the presence or absence of career concerns.

Our model of decision-making in the FOMC is a model of voting with interdependent values augmented by career concerns. The baseline follows Iaryczower and Shum (2012): agents have state-contingent payoffs, which depend on how their policy recommendation (*hawkish* or *dovish*) matches whether the optimal monetary policy is expansionary (in a “recession”) or contractionary (in a “boom”). Members differ in the disutility of implementing an inappropriate policy (expansion in booms or contraction in recessions). Each member receives a private signal about the state, drawn from a normal distribution with a state-dependent mean and an individual-specific variance, with the signal precision capturing a member’s competence. To incorporate career concerns in a tractable way while retaining heterogeneity across members, we assume that each member is either an *expert*, who perfectly observes the state and votes accordingly, or a *strategic type*, who is imperfectly informed and cares about outsiders’ beliefs that they are experts. We characterize equilibrium behavior under two informational regimes. In the Transparent regime, outsiders observe all individual recommendations; in the Opaque regime, they observe only the aggregate outcome.

We show that the model’s parameters are point identified. Two elements drive our identification strategy. First, following Hotz and Miller (1993), Bajari, Benkard, and Levin (2007), and Iaryczower and Shum (2012), we recover conditional choice probabilities, decompose the game into individual decision problems, and identify agent-level parameters from observed behavior and best responses (which take the form of cutoff strategies). Within each regime, conditional choice probabilities are identified using standard mixture-model arguments (e.g., Allman, Matias, and Rhodes (2009)). These probabilities allow us to recover equilibrium cutoffs, individual competence, and the probability of being an expert type, leveraging the monotonic mapping between choice probabilities and equilibrium cutoffs. Second, we exploit the regime change to separately identify policy preferences and career concerns by comparing the equilibrium conditions across the two regimes.

Consistent with previous research on the FOMC, we find that the typical Governor is more dovish and better informed about the state of the economy than the typical President. Unlike prior studies, however, we are also able to estimate the value that policymakers place on their reputation. We find that, in addition to being more

dovish and better informed, the typical Governor assigns greater weight to reputation than the typical President. That said, there is substantial heterogeneity within both groups in terms of preferences, information, and career concerns: over the sample period, 24% of Presidents are more dovish than the median Governor, 16% are more competent, and 16% place a higher value on reputation than the median Governor.

With the estimates in hand, we turn to our main objective: quantifying the effect of career concerns on outcomes. Our key outcome of interest is the probability that a member issues a correct recommendation in each state of the economy (booms and recessions), which we compute from the estimated model. We also examine the heterogeneous effects of conformity using potential outcomes implied by the structural framework. We obtain five main results:

First, to assess how career concerns shape behavior within a given institutional regime, we compare outcomes *with* career concerns- as observed in the data-to those that would arise in a counterfactual scenario *without* career concerns. We find that, in both regimes, career concerns tend to increase the median error in booms and reduce it in recessions. In the Opaque regime, however, their effect on the probability of a correct recommendation is relatively small, whereas under Transparency the impact is substantial. Given the higher frequency of recessionary states in the data, we show that career concerns have a moderately positive effect on the *ex ante* probability of a correct recommendation.

Second, we compare informational regimes by evaluating a counterfactual in which each regime is held fixed throughout the sample. We find that Transparency generally outperforms Opacity during recessions, whereas the Opaque regime dominates in booms. When considering the *ex ante* probability of a correct recommendation, however, Transparency consistently yields higher performance across the full sample. In other words, for an unbiased Principal-who values policy errors in recessions and booms equally-Transparency leads to better outcomes throughout the sample period.

Third, we focus on the mechanisms behind the effect of changes in transparency, by quantifying the informational distortions induced by career concerns in both Opaque and Transparent regimes. We show that committee members – particularly Presidents – engage in anti-pandering, contradicting the Principal’s prior beliefs, and that transparency amplifies this tendency. Quantitatively, however, the dominant behavioral change is due to *mimicking*: agents suppress their policy biases to appear competent

and unbiased. This force generally outweighs the distortions from pandering.

Fourth, we quantify the *strategically induced* component of conformity, separating it from correlation in actions driven by affiliated values. We find that the effects of career concerns on conformity are highly heterogeneous across meetings. Consistent with the reduction in disagreement documented in reduced-form results, we estimate that strategic conformity is close to zero in the Opaque regime, but shifts sharply on average with the move to Transparency—it increases for Governors and declines for Presidents. Using our structural estimates, we then recover heterogeneous treatment effects by replicating the key difference-in-differences coefficient in terms of potential outcomes, computed for each meeting under counterfactual conditions. The resulting average treatment effect closely matches the reduced-form point estimate, while also revealing substantial heterogeneity across policy meetings.

Fifth, our reduced-form analysis shows that, while Transparency reduces dissent on average, Presidents become more responsive to regional shocks under Transparency than under Opacity. We revisit this finding using the structural model to examine how regional shocks affect *equilibrium* outcomes. We find that greater heterogeneity in regional shocks has no effect under Opacity, but increases policy errors and reduces strategic conformity under Transparency. These results highlight how transparency amplifies both the informational and strategic impact of local economic conditions, particularly as they shape the relative accuracy of Presidents compared to Governors.

The rest of the paper is organized as follows. Section 2 lays out the connection to the literature. In Section 3 we discuss the institutional context and provide preliminary reduced-form evidence. In Section 4, we present the model, and characterize best responses under both regimes. Section 5 discusses identification and estimation. In Section 6 we provide our results. We start by describing structural estimates, and then address our main questions in Section 6.1. We offer our concluding thoughts in Section 7.

2 Related Literature

To the best of our knowledge, our paper is the first to explicitly incorporate career concerns in an empirical equilibrium model of decision-making in committees. In doing this, we build on both the theoretical and empirical literatures analyzing the

effect of career concerns on behavior, with particular emphasis on studies of career concerns within the FOMC and other central banks.

Theory. In our model, agents are concerned about their reputation for being competent and unbiased. The first model studying reputation for competence is provided by Holmström (1982) (Section 3.2). Since then, a large theoretical literature studied the distortions in behavior induced by career concerns.² In the single decision-maker case, Canes-Wrone, Herron, and Shotts (2001), Levy (2004), Prat (2005), and Ottaviani and Sørensen (2006) show that depending on whether the agent knows her type or not, reputation for competence can create an incentive to either favor or contradict the action that the principal believes to be optimal given her prior beliefs (i.e., to “pander” or “anti-pander” to the principal).³

The idea that career-concerned agents may disregard their private signals to act as a competent agent is expected to behave *a priori* predates the single-agent setting. In a model with multiple career-concerned advisors, Scharfstein and Stein (1990) establish conditions for a *herding* equilibrium, in which the second advisor mimics the first regardless of her private (binary) signal (see also Zwiebel (1995); Prendergast and Stole (1996); Effinger and Polborn (2001); Ottaviani and Sørensen (2001)). In such environments, advisors’ actions convey information about the state, altering the principal’s prior expectations about the optimal decision. A similar information externality arises in committees, even with simultaneous voting. Here, “herding” takes the form of *conformity*, i.e., incentives to present a unified front (see Levy (2007); Visser and Swank (2007)). The logic is straightforward: with interdependent values, higher competence increases the correlation between a member’s vote and the state, thereby strengthening the correlation across members’ votes. Because disagreement signals a lack of competence, career-concerned agents have incentives to conform. Importantly, Levy (2007) shows that reputation concerns can induce strategic voting

²We distinguish this literature from agency models studying moral hazard, as in the basic Holmström (1982) model. In the FOMC context, the “effort” dimension has been interpreted as effort in information acquisition about economic conditions (see Hansen, McMahon, and Prat (2018)). Following Blinder (2008), we believe that given the robust staff working for both Presidents and Governors, and the importance of monetary policy decisions, the adverse selection problem is first order in this setting.

³Maskin and Tirole (2004) show that pandering can also emerge in equilibrium if the decision-maker wants to signal to the principal that she has “congruent” preferences. In a similar context, Morris (2001) and Stasavage (2007) show that with sufficiently high reputational concerns, the incentives to pander prevent transmission of information in equilibrium.

even under secretive procedures.

In order to focus on informational distortions from career concerns, the above models generally assume homogeneous preferences. In contrast, our model allows for heterogeneity in policy preferences among committee members. In this richer environment, signaling for reputation can also lead to *mimicking*, as in Finan and Mazzocco (2021). We show that mimicking can dominate the distortions induced by anti-pandering and conformity, leading to welfare improvements.

Transparency in the FOMC. Meade and Stasavage (2008) are the first to examine the effects of changes in FOMC transparency, focusing on the average probability of dissent between members’ policy views and the Chairman’s proposals before and after the 1993 “transparency shock” (see also Swank, Visser, and Swank (2008)).⁴ They show that transparency reduces dissent with the Chair, and interpret this as evidence of conformity. Hansen, McMahon, and Prat (2018) examine the same transparency shock more broadly, applying topic models to FOMC transcripts. They show that, following the transparency shock, deliberations became more quantitatively oriented and focused on economic issues. Using a DiD specification – with Federal Reserve experience as a proxy for career concerns – they find that less experienced members spoke more like Greenspan than did their more experienced counterparts. They interpret this finding to say that stronger career concerns increase herding.⁵

Whereas previous studies focus primarily on dissent as the key outcome of deliberation, our approach adopts a broader perspective. We measure efficiency by the state-contingent probability of a correct recommendation conditional on the state of the economy, which is directly tied to welfare.⁶ To our knowledge, we are also the first

⁴Meade and Stasavage (2008) analyze seventy-two FOMC meetings held between 1989 and 1997 under Chairman Greenspan. Chappell Jr, Havrilesky, and McGregor (1997) were the first to study internal discussions of the FOMC (rather than formal votes), focusing on Chairman Burns. Meade (2005) and Chappell, McGregor, and Vermilyea (2005) extend this analysis to the Greenspan era, showing that disagreement during deliberations was substantially greater than the low rate of official dissent recorded under Greenspan.

⁵To compare the “positive” effect of greater effort with the “negative” effect of stronger anti-herding, the authors compute an influence score in the spirit of PageRank. They find that inexperienced members became more influential in shaping colleagues’ topic coverage after transparency. This suggests that for members with strong career concerns, the effort effect dominates the informational loss from anti-herding.

⁶Papers adopting a similar approach in other contexts include Iaryczower and Shum (2012), Iaryczower, Lewis, and Shum (2013), Hansen, McMahon, and Rivera (2014), López-Moctezuma (2016), Iaryczower and Katz (2016), Iaryczower, Shi, and Shum (2018), and Canen and Iaryczower (2024). López-Moctezuma (2016) is the first to apply this framework to the FOMC.

to quantify distortions from pandering, conformity, and signaling within a structural framework.

Other structural work on Central Banks. Hansen and McMahon (2016) estimate a structural model of monetary policymaking in the Bank of England’s Monetary Policy Committee (MPC). Their central objective is to test whether policymakers behave more aggressively against inflation early in their careers but adopt looser policies as their tenure progresses. Unlike our framework, committee members in their model do not face career concerns. Instead, rookie members signal hawkish preferences to the public to influence inflation expectations, following the mechanism of Sibert (2003) (see also Sibert (2009)).⁷ Hansen, McMahon, and Rivera (2014) study information aggregation in the MPC, building on Iaryczower and Shum (2012), while López-Moctezuma (2016) analyze sequential deliberation in the FOMC, emphasizing the role of social learning in monetary policymaking. None of these papers incorporate career concerns.

Career Concerns More Broadly. Our paper relates to a series of studies that use structural estimation to analyze how elected politicians respond to electoral incentives. Diermeier, Keane, and Merlo (2005) and Lim (2013) estimate dynamic models of congressional careers and of elected versus appointed judges, respectively. Iaryczower, Lopez-Moctezuma, and Meirowitz (2024) estimate a dynamic model in which senators choose policy positions and campaign advertising in response to their polling performance. Their focus on within-cycle dynamics allows them to estimate the value of career concerns at the individual level. Unlike our approach, however, these studies do not quantify the informational distortions induced by career concerns. Sieg and Yoon (2017) and Aruoba, Drazen, and Vlaicu (2018) structurally estimate models in which a single officeholder takes an action while in government and may seek reelection. As in Banks and Sundaram (1998), there is moral hazard and adverse selection, but no joint learning about the economy and the politician’s type, and thus no scope for distortions such as pandering.⁸ In these models, reelection

⁷For other studies on the Bank of England’s MPC, see Gerlach-Kristen (2004), Harris and Spencer (2009), Hix, Høyland, and Vivyan (2010), and Hansen, McMahon, and Rivera (2014).

⁸In Aruoba, Drazen, and Vlaicu (2018), effort stochastically generates good outcomes for voters, but types differ in the cost of effort. In Sieg and Yoon (2017), voters are uncertain about the politician’s ideal point and update beliefs based on observed policy.

incentives generate positive selection and encourage effort or policy moderation, but strategic interactions between agents are absent, unlike in the FOMC context.

3 Context, Data and Facts

Institutional Context. The Federal Open Market Committee (FOMC) is the body responsible for implementing U.S. monetary policy, primarily by targeting the federal funds rate (FFR).⁹ The FOMC consists of the seven members of the Board of *Governors* of the Federal Reserve System and the twelve Regional Federal Reserve Bank *Presidents*. Governors are appointed by the U.S. president, with Senate confirmation, to (staggered) fourteen-year terms. Fed Presidents are appointed by their Regional Bank’s Board of Directors—nine members representing banking, agricultural, commercial, industrial, and public interests within the district—subject to approval by the Board of Governors, and serve renewable five-year terms. At any given meeting, twelve members hold voting rights: the seven Governors, the President of the Federal Reserve Bank of New York, and four of the remaining eleven district Presidents, who rotate annually as voting members. The remaining non-voting Presidents still attend FOMC meetings, participate in economic assessments, and provide policy recommendations.

The FOMC meets about eight times per year, following a standard protocol with two stages: *(i)* the economic outlook discussion and *(ii)* the policy discussion. In the economic outlook discussion, the Board’s staff presents an overview of current economic conditions and forecasts of key macroeconomic variables. Members then respond with questions, comments, and their own assessments. In the policy discussion, the staff outlines potential policy alternatives, after which members state and debate their preferred policy positions. Finally, the Chair proposes a policy directive, which must be approved by a majority of voting members.

Formal votes in the FOMC are largely consensus-driven, with rare instances of dissent. In fact, *every* adopted policy aligns with the Chair’s directive, which in turn generally reflects the majority view expressed during the policy discussion (see Table A in the Appendix). In our baseline model, we therefore assume that policy outcomes follow simple majority rule over voiced preferences and focus on members’

⁹From 1979 to 1989, the FOMC replaced FFR targeting with non-borrowed reserves (1979-1982) and borrowed reserves (1982-1989) in an effort to more closely control monetary aggregates.

recommendations during the discussion stage.¹⁰ The near-unanimity of official votes contrasts sharply with the higher level of disagreement in policy discussions: between 1970 and 2008, 22% of voiced recommendations diverged from the final decision, compared to only 7% dissent in formal votes.

Data. Our sample covers all available FOMC meetings between 1970 and 2008, ending in the midst of the "Great Recession," when the FOMC shifted its primary policy instrument to quantitative easing. In total, we analyze 347 meetings, involving 102 unique FOMC members and 5,570 individual policy recommendations. Whenever possible, we collect members' desired federal funds rate (FFR) targets directly from their statements during the policy discussion; otherwise, we infer their recommendations based on contextual information, following the coding standards established by Chappell, McGregor, and Vermilyea (2005) and Meade (2005).¹¹ For estimation of the structural model, we construct a binary outcome variable for each member indicating whether the member supported a "hawkish" or "dovish" policy alternative in each meeting.¹²

We complement the policy data with information on macroeconomic fundamentals at the time of each meeting, and member's individual characteristics. To control for members' expectations of the economic environment, we include three-month-ahead forecasts of inflation, unemployment, and real GDP growth prepared by the Board of Governors' staff. These forecasts, distributed to members before each meeting, are available in the *Green Book*. To capture inertia in policymaking, we include the FFR from the week prior to each meeting. Because monetary aggregates have at times served as intermediate policy targets, we also control for recent money growth, measured as the mean of the last three available monthly figures prior to each meeting. In terms of individual characteristics, we include gender, career background, FOMC

¹⁰In Appendix F, we extend our model to consider formal votes. We show that if we include a cost of dissenting in formal votes, our model – evaluated at our main structural parameter estimates – explains variation in formal votes well. We conclude that our simplified representation of the Opaque regime does not introduce a significant bias in our estimates.

¹¹Between 1979 and 1982, the FOMC's primary instrument shifted from the FFR to non-borrowed reserves. To code recommendations in this period, we use the mapping between borrowing targets and FFR targets from the *Blue Book*. When mappings are given as ranges, we use their midpoint as the voiced preference. Results are robust to excluding these meetings (see Table B3 in the Appendix).

¹²In 86% of cases, members consider at most two staff proposals. When more than two are available, we define a hawkish action as supporting an FFR equal to or higher than the staff median proposal. Figure A shows the relationship between individual FFR recommendations and the *Blue Book* policy alternatives.

experience, institutional affiliation, and, for Governors, the partisan affiliation of the appointing President.¹³ For Fed Presidents, we also capture local economic conditions using the regional unemployment gap at each meeting.¹⁴

Figure 1 plots the evolution of FOMC policy recommendations and their dispersion over time, alongside FOMC staff forecasts for inflation (left-hand panel) and unemployment (right-hand panel). The dynamics of FOMC members’ recommendations closely align with the committee’s aggregate expectations of the economic environment, with the FFR generally exhibiting a positive correlation with expected inflation and inversely related to expected unemployment. This relationship underscores the crucial role of macroeconomic expectations in shaping FOMC members’ policy preferences, which we account for in our empirical analyses below.

Table A in the Appendix presents summary statistics on the distribution of policy recommendations and the average macroeconomic conditions across different Chairmanships in the sample. The Burns and Miller regimes were marked by rising inflation and a significant slowdown in economic growth. During the early years of Volcker’s tenure, the FOMC implemented an aggressive monetary tightening cycle to combat the high inflation of the 1970s. While this policy response led to the economic recession of the early 1980s, it also contributed to increased dispersion in policy recommendations among FOMC members. Greenspan’s tenure coincided with sustained economic growth, well-anchored inflation, and a stable FFR, alongside reduced dispersion in policy recommendations. However, the prosperity of the 1990s gave way to the most severe economic crisis since the Great Depression during Bernanke’s tenure, prompting the committee to aggressively lower the FFR to the zero lower bound.

Transparency Shock In order to provide a model-free assessment of the effect of Transparency on outcomes, we use a natural experiment that abruptly changed FOMC communication with the public, following Meade and Stasavage (2008) and

¹³Following Adolph (2013), we measure career background as the fraction of a member’s career spent in seven mutually exclusive categories up to their most recent appointment to the FOMC: (i) financial (private banking), (ii) government (non-Fed, non-Treasury bureaucrats), (iii) finance ministry (Treasury Department), (iv) central bank (Federal Reserve staff), (v) economics (academic economists), (vi) business (private sector excluding banks), and (vii) other (e.g., international organizations). Experience is measured with an indicator for rookie status (tenure below the 25th percentile in the sample).

¹⁴Monthly regional unemployment by Federal Reserve district is available from 1990 onward. We extend the series back to 1976 by computing population-weighted averages using county and state unemployment rates, county population, and Federal Reserve district boundaries.




Figure 1: Policy Recommendations and Economic Environment. The left panel shows the maximum and minimum FFR recommendations by FOMC meeting along with FOMC’s staff three-month ahead inflation forecast. The right panel shows the maximum and minimum FFR recommendations by FOMC meeting along with FOMC’s staff unemployment forecast.

Hansen, McMahon, and Prat (2018). In October 1993, Fed Chairman Alan Greenspan acknowledged the existence of unedited FOMC meeting transcripts going back to 1976, which were used in the process of preparing meeting minutes. Up until that point, FOMC members were unaware that meeting transcripts had been preserved (Lindsey (2003)).¹⁵ After sustained pressure from Congress, the FOMC decided to release past meeting transcripts with a lag of five years at the meeting of November of 1993, and in February 1994, the FOMC decided to release all future transcripts. The measure was not seen as a minor procedural change within the FOMC. As Greenspan himself acknowledged (extracted from Lindsey (2003)),

“What does bother me is putting on the table and in the public our deliberations about potential future [actions] and our inclinations and the conditions that would drive [such actions]. Interacting with the market would then alter our decision-making process and create a significant loss of flexibility in the actions we might have undertaken.”

This shock to the transparency of FOMC deliberations means that we can analyze *individual* policy recommendations both under the *Opaque* regime, where meeting

¹⁵Only Governor Mullins and President Boehne had become aware in previous months of the existence of the transcripts. President Melzer later testified learning of the transcripts’ existence in 1989.

participants deliberated monetary policy under the assumption that their discussions were private (prior to November 1993), and under a *Transparent* regime where FOMC members deliberated knowing their statements would be public record within five years.

In our reduced-form analysis we focus on conformity/disagreement within the FOMC. Later on, we use our structural model estimates to revisit these results. Following Meade and Stasavage (2008) and Swank, Visser, and Swank (2008), we measure disagreement within the FOMC using a binary indicator based on the gap between each member’s policy recommendation and the Chair’s directive.¹⁶ To capture systematic differences in the *direction* of disagreement, we also compute the average difference between a member’s recommended FFR and the implemented FFR. Following Hansen, McMahon, and Prat (2018), we also construct a broader measure of disagreement based on members’ statements during deliberations, which captures the dissimilarity between the *topics* covered by each member and those covered by the Chair.¹⁷

Key Facts. In this section, we exploit the transparency shock and the dual membership in the FOMC to document three basic facts: relative to behavior in the Opaque regime, Transparency

1. reduces the “hawkishness gap” between Presidents and Governors,
2. reduces Governors’ disagreement rate relative to Presidents,
3. increases Presidents’ sensitivity to regional shocks.

To examine the effect of the transparency shock on hawkishness and disagreement, we exploit differences in outcomes between Governors and Presidents before and after the shock. This approach allows us to isolate reputational incentives across the two

¹⁶We classify a recommendation as disagreeing with the Chair when the member’s FFR proposal deviates by at least 25 basis points from the Chair’s directive. Alternatively, we define disagreement using our binary “hawkish vs. dovish” classification of recommendations relative to the Chair. Our reduced-form results are robust to using either definition (see Table B6 in the Appendix).

¹⁷We apply a probabilistic topic model (Roberts, Stewart, Tingley, Airoldi, et al. (2013)) to recover the prevalence of policy topics in each member’s speech during the policy discussion, and then compute the Kullback-Leibler (KL) divergence between the resulting topic distributions of each member and the Chair. For vocabulary and model selection, we follow Hansen, McMahon, and Prat (2018). See Appendix C.

groups while controlling for economic fundamentals and FOMC member characteristics. Specifically, we estimate the following regression:

$$y_{it} = \alpha_{a[i]} + \delta_t + \beta T_t \times Pres_i + \mathbf{X}_{it}'\gamma + \epsilon_{it}, \quad (3.1)$$

where the dependent variable y_{it} represents one of our three measures of disagreement between a member and the Chairperson. The term $\alpha_{a[i]}$ captures affiliation fixed effects (Republican-appointed Governors, Democratic-appointed Governors, and the twelve regional bank Presidents), while T_t is a dummy variable for the transparency regime (equal to one after November 1993 and zero otherwise). The variable $Pres_i$ indicates whether the member is a Fed President, and $\mathbf{X}_{it}$ is a vector of covariates including individual characteristics such as gender, experience, and career background. The independent effect of $Pres_i$ is absorbed by affiliation fixed effects, while the transparency shock (T_t) and macroeconomic conditions are captured by time-fixed effects. All results are reported with Driscoll and Kraay (1998) standard errors, which are robust to cross-sectional dependence, serial correlation, and heteroskedasticity.

To compare our results with the previous literature, we also estimate a “before and after” specification, in which we remove the time FEs and include the transparency dummy (T_t), controlling for economic fundamentals, including staff forecasts for real GDP, inflation, and unemployment. This alternative specification allows us to consider the level change induced by the switch to the Transparent regime. The obvious caveat is that the before and after specification only imperfectly controls for changing economic fundamentals.

Table 1 presents the main estimates for our three measures of disagreement (see Table Table B1 for the full results).¹⁸ Columns (1) and (2) report the results for disagreement in policy recommendations. Column (1) presents the “before and after” specification, controlling for macroeconomic fundamentals, while Column (2) presents the results of the DiD specification. The results from the “before and after”

¹⁸Our baseline results in Table 1 include all observations in the sample. In Table B3, we show that our findings remain consistent even when excluding meetings from 1993, when at least three FOMC members were aware that transcripts existed and that their statements might eventually become public. Our results are also robust to excluding the period 1979-1992, when the FOMC primarily targeted non-borrowed reserves. As a further robustness check, we split the Greenspan sample into pre- and post-transparency shock periods and conduct a placebo test, using the midpoint dates as the placebo transparency changes (October 1990 for the *Opaque* period, and December 1999 for the *Transparent* period). We present the results in Table B4.

	Disagreement					
	FFR (Rec.)		FFR (Direction)		Topic	
	(1)	(2)	(3)	(4)	(5)	(6)
$T_t \times Pres_i$	0.072** (0.028)	0.059** (0.028)	-0.062** (0.026)	-0.066** (0.028)	0.976*** (0.296)	1.040*** (0.260)
Transparency	-0.096** (0.047)		0.046 (0.037)		-1.156* (0.682)	
Mean (Govs., Opaque)	0.29	0.29	-0.02	-0.02	5.69	5.69
Affiliation FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	No	Yes	No	Yes
Individual Chars.	Yes	Yes	Yes	Yes	Yes	Yes
Economic Chars.	Yes	No	Yes	No	Yes	No
Observations	5,090	5,090	5,090	5,090	3,913	3,913
R ²	0.137	0.026	0.015	0.023	0.163	0.193

Note: *p<0.1; **p<0.05; ***p<0.01

Driscoll-Kraay Standard Errors in Parentheses

Table 1: Disagreement by Transparency and Affiliation

specification show that transparency is associated with a significant reduction of disagreement on average, consistent with the result of Meade and Stasavage (2008) for the Greenspan era (1989-1997). The DiD coefficient, moreover, shows that the reduction in disagreement for the average Governor (0.10) more than doubles the reduction in disagreement for the average President (0.04). Columns (5) and (6) reproduce the specifications of (1) and (2), measuring disagreement using the textual content of FOMC members' statements during policy deliberation, as in Hansen, McMahon, and Prat (2018).¹⁹ The results are consistent with the findings for the disagreement in policy recommendations. Columns (3) and (4) present the results for the direction of disagreement. As shown in Figure B1, and more formally in Table B2, Fed Presidents exhibit higher rates of disagreement with the policy directive than Governors, and consistently advocate higher FFR targets. The results in Columns (3) and (4) show that in the Transparent regime, Presidents shift toward a more dovish stance, while Governors become more hawkish, reducing the "hawkishness gap" between the two groups.

We now turn to our third stated fact, that Transparency increases Presidents' sensitivity to regional shocks. While both Governors and Presidents consider regional

¹⁹In Hansen, McMahon, and Prat (2018), the specification doesn't compare the behavior of Presidents and Governors to the shock, but instead uses a measure of experience in the Federal Reserve to proxy for career concerns.

macroeconomic disparities, Presidents are particularly attuned to local economic conditions in their respective districts.²⁰ Fed Presidents’ policy recommendations (Chappell Jr, McGregor, and Vermilyea (2008)), voting behavior (Gildea (1992); Meade and Sheets (2005)) and their public communications (Hayo and Neuenkirch (2013)) have been shown to respond to regional economic conditions. This raises the question of whether their attention to local interests would be heightened or lessened with the higher scrutiny brought by Transparency. If Presidents’ career concerns were solely driven by signaling competence to a national audience, Transparency might incentivize a more balanced focus on national conditions. If instead Presidents partly serve local interests, one might expect that greater transparency would increase their sensitivity to regional shocks when providing their policy recommendations.

To examine how Fed Presidents respond to regional economic shocks under different transparency regimes, we estimate the following regression for the subset of FOMC members serving as Presidents:

$$y_{it} = \alpha_{D[i]}^{reg} + \delta_t^{reg} + \beta^{reg} T_t \times unemp_{D[i],t} + \mathbf{X}_{it}' \gamma^{reg} + \epsilon_{it}, \quad (3.2)$$

where $unemp_{D[i],t}$ denotes the difference between the unemployment rate in District D and national unemployment at meeting t . (National economic conditions, including national unemployment, are absorbed by the time fixed effects, δ_t^{reg}). We estimate the model for our three measures of disagreement: the probability of disagreement in recommendations, the direction of disagreement (i.e., the difference between the desired rate and the policy directive), and the textual disagreement measure.

We present the results in Figure 2 (see Table B5 in the Appendix). The estimate of β^{reg} is negative and statistically significant across all outcomes only in the Transparent regime. Under Transparency, Fed Presidents are more inclined to support easier monetary policies when regional unemployment rises relative to the national level, leading to reduced disagreement with the Chair’s policy recommendations. Similarly, Presidents’ speech during policy deliberations increasingly aligns with that of the Chair when unemployment in their region rises relative to the national average. Figure 2 illustrates the magnitude of these effects by simulating disagreement outcomes for the

²⁰Fed Presidents share their assessments of local economic conditions to complement the aggregate data compiled by the Board of Governors’ staff. These insights often include confidential information on local businesses, markets, and key industries, and can help identify emerging economic trends before they appear in national-level indicators (Goodfriend (1999)).




Figure 2: Fed Presidents' Disagreement by Regional Unemployment and Regime. Solid lines represent simulated outcomes as we move $regunemp_{it}$ from the minimum to the maximum in the data. Shaded regions cover 95% confidence intervals.

median Fed President as regional unemployment shifts from its minimum to maximum value in the data. The estimated probability of disagreement drops from 0.29 to zero when regional unemployment shifts from -3% to +3.35% above the national average. This increase in the regional-national unemployment gap leads the median Fed President to adjust their policy recommendation from 14 basis points above the adopted FFR to 14 basis points below it. The same shift in relative unemployment reduces topic disagreement by more than 30% in the Transparent regime.

4 The Model

We begin by describing the basic decision-making model (preferences and information), without reputation considerations. We then augment the game with career concerns.

Basic Decision-Making Environment. A committee of n members makes a binary policy decision $d \in \{0, 1\}$, of whether to implement a dovish ($d = 0$) or a hawkish ($d = 1$) monetary policy. Committee members make decisions under two different regimes. In the Opaque regime, individual committee recommendations $v_i \in \{0, 1\}$ are secret, and only the final decision $y \in \{0, 1\}$ is observed. In the Transparent regime, deliberations are observed by the public. We let $\mathcal{Y}^r$ denote the outcome observed by the public in regime $r = O, T$.

The economy can be in a recessionary or inflationary state, $\omega \in \{0, 1\}$. Committee members have state-contingent preferences, $u_i : \{0, 1\}^2 \rightarrow \mathbb{R}$, where $u_i(v_i, \omega)$ denotes i 's payoff from recommending $v_i \in \{0, 1\}$ in state $\omega \in \{0, 1\}$. Agent i gets a policy payoff $u_i(1, 1) = u_i(0, 0) = 0$ when their decision matches the state, a payoff $u_i(1, 0) = -\pi_i$ when she recommends a hawkish policy in a recession, and a payoff $u_i(0, 1) = -(1 - \pi_i)$ when she recommends a dovish policy in a boom. This specification implies that, absent reputational considerations, i favors a hawkish policy ($d = 1$) when her posterior belief that the economy is in a boom exceeds the threshold π_i . The parameter $\pi_i \in (0, 1)$ thus captures the degree of i 's dovishness. Committee members are only partially informed about the realization of the state, and have prior beliefs $\rho \equiv \Pr(\omega = 1)$ that the economy is in an inflationary state. In particular, agent i observes a private signal $s_i = \omega + (1/\theta_i)\varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, 1)$, and $\theta_i \in \mathbb{R}_+$ parametrizes the precision of i 's information (i.e., competence).

Career Concerns. Our aim is to capture career concerns in a tractable manner, while preserving a rich space of characteristics for each committee member. To do this, we introduce an *expert* behavioral type, while allowing the policy preferences, value of reputation, and quality of information of the *strategic* type to vary flexibly across members.

Specifically, we assume that with probability $p_i \in (0, 1)$, member i is an unbiased *expert* ($\tau_i = e$), who observes the realization of the state, and makes a recommendation $v_i = \omega$. With probability $1 - p_i$, member i is a strategic decision-maker with the preferences, information and payoffs described above. The characteristics of a strategic DM i are common knowledge for both committee members and principals, and can differ across committee members.

Strategic DMs care about their reputation for competence, which is captured by the Principals' beliefs that the agent is an expert, $P_i^r(\mathcal{Y}^r) \equiv \Pr(\tau_i = e | \mathcal{Y}^r)$. Formally, agent i 's expected utility of recommending $v_i \in \{0, 1\}$ in regime r when observing signal s_i given strategy profile β^r is

$$U_i^r(v_i | s_i) = E[u_i(v_i, \omega) | s_i] + \Delta_i E_{v_{-i}}[P_i^r(\mathcal{Y}^r(v_i, v_{-i}); \beta^r) | s_i], \quad (4.1)$$

where $\Delta_i \in \mathbb{R}_+$ captures how much the agent cares about her reputation for competence relative to the decision at hand, and β^r denotes the strategy profile in regime

r (defined formally below).

To capture the possibility that regional Presidents have different career incentives, we allow heterogeneous priors ρ_i by the principals. In particular, we leverage the reduced-form evidence on the regional bias of Presidents' behavior and assume that President i 's principal forms biased beliefs about ω following local economic conditions. Letting x_i denote regional unemployment over the national level, we assume that i 's principal updates her beliefs about ω assuming that $x_i \sim \mathcal{N}(1 - 2\omega, 1/\zeta^2)$. With this,

$$\frac{\Pr(x_i|\omega = 0)}{\Pr(x_i|\omega = 1)} = \exp(2\zeta^2 x_i) \Rightarrow \rho_i \equiv \Pr(\omega = 1|x_i) = \left[1 + \left(\frac{1 - \rho}{\rho}\right) \exp(2\zeta^2 x_i)\right]^{-1},$$

which boils down to the common prior case for $\zeta \rightarrow 0$.

Let $\beta_i^r(\cdot) : \mathbb{R} \rightarrow [0, 1]$, where $\beta_i^r(s_i) = \Pr(v_i = 1|s_i; r)$ denotes the probability that a strategic DM i recommends a hawkish monetary policy, $v_i = 1$ when receiving signal s_i in regime $r = O, T$. Since experts are behavioral types, we refer for convenience to $\beta_i^r(\cdot)$ as agent i 's strategy in regime r . An equilibrium in regime $r \in \{O, T\}$ is a Perfect Bayesian equilibrium (PBE) of the corresponding game. Since the agents' payoffs only depend on the Principal's beliefs, an equilibrium in regime $r \in \{O, T\}$ can be equivalently defined as a Bayesian Nash equilibrium (BNE) of the corresponding game among the agents, with the added constraint that the principal's beliefs are determined by Bayes' rule from agents' equilibrium strategies.

4.1 Preliminaries

In our estimation strategy, we follow an approach inspired by Hotz and Miller (1993) and Bajari, Benkard, and Levin (2007) for dynamic games, and Iaryczower and Shum (2012) for voting with interdependent values. In this approach, we first estimate conditional choice probabilities, and then use equilibrium conditions to obtain an estimate of the model parameters. In our context, these conditional choice probabilities are given by:

$$\gamma_{i,\omega}^r \equiv \Pr(v_i = 1|\omega; r) = p_i\omega + (1 - p_i)\sigma_i^r(\omega) \quad (4.2)$$

for regime $r \in \{O, T\}$, individual $i \in I$, and state $\omega \in \{0, 1\}$. Here

$$\sigma_i^r(\omega) \equiv \Pr(v_i = 1 | \omega, \tau_i = a; r) = \int_s \beta_i^r(s) f(s | \theta_i, \omega) ds$$

is the conditional probability that a strategic type of member i recommends a hawkish policy in state $\omega \in \{0, 1\}$ and regime $r \in \{O, T\}$. A crucial feature of this approach in our context is that it allows us to decompose the game into n independent decision problems, enabling the identification of individual parameters from best responses, given the observed data. With this in mind, the next section focuses on characterizing best responses in both the Opaque and Transparent regimes, taking the conditional choice probabilities as given.

4.2 Best Responses in the Transparent and Opaque Regimes

Recall that from (4.1), agent i 's expected utility of recommending $v_i \in \{0, 1\}$ in regime $r \in \{O, T\}$ when observing signal s_i given strategy profile β^r is

$$U_i^r(v_i | s_i) = E[u_i(v_i, \omega) | s_i] + \Delta_i E_{v_{-i}}[P_i^r(\mathcal{Y}^r(v_i, v_{-i}); \beta^r) | s_i],$$

Noting that $E[u_i(1, \omega) - u_i(0, \omega) | s_i] = \Pr(\omega = 1 | s_i) - \pi_i$, recommending a hawkish monetary policy is a best response after observing a signal s_i , iff

$$\underbrace{\Pr(\omega = 1 | s_i) \geq \pi_i}_{\text{No CC}} + \underbrace{\Delta_i E_{v_{-i}}[P_i^r(\mathcal{Y}^r(0, v_{-i}); \beta^r) - P_i^r(\mathcal{Y}^r(1, v_{-i}); \beta^r) | s_i]}_{\text{reputation}} \quad (4.3)$$

In the absence of career concerns, i recommends a hawkish policy iff $\Pr(\omega = 1 | s_i) \geq \pi_i$. Since i 's belief that the economy is in a boom is increasing in s_i , it follows immediately that in the equilibrium of either regime, if i has no career concerns, her best response is a cutoff strategy; i.e., $\beta_i^r(s_i) = 0$ for all $s_i < \xi_i^{noCC}$ and $\beta_i^r(s_i) = 1$ for all $s_i > \xi_i^{noCC}$, where $\xi_i^{noCC} \in \mathbb{R}$ is given by $\Pr(\omega = 1 | \xi_i^{noCC}) \equiv \pi_i$. As a result, $\sigma_i^r(\omega) = 1 - \Phi([\xi_i^{noCC} - \omega] \theta_i)$.

The second term on the right hand side of (4.3) captures career concerns; the net expected reputational gain of giving a dovish instead of a hawkish recommendation given a signal s_i . Since i has more information than the Principal, the uncertainty doesn't arise from the Principal's beliefs about i given $\vec{v}$, but from the randomness

in other committee members' recommendations. In turn, their expected behavior is fully captured by the conditional choice probabilities in each state, $\gamma_{-i,\omega}^r$. Thus, we can write (4.3) as

$$\beta_i^r(s_i) = 1 \Leftrightarrow \Pr(\omega = 1|s_i) \geq \pi_i + \Delta_i E_\omega [\Psi_i^T(\omega; \beta_i^r, \gamma_{-i,\omega}^r)|s_i] \quad (4.4)$$

where

$$\Psi_i^r(\omega; \beta_i^r, \gamma_{-i,\omega}^r) \equiv E_{v_{-i}} [P_i^r(\mathcal{Y}^r(0, v_{-i}); \beta_i^r, \gamma_{-i,\omega}^r) - P_i^r(\mathcal{Y}^r(1, v_{-i}); \beta_i^r, \gamma_{-i,\omega}^r)|\omega]$$

is i 's expected reputational gain from recommending a dovish vs a hawkish policy in regime r conditional on the state $\omega \in \{0, 1\}$. The difference in incentives across regimes therefore operates exclusively through the term $\Psi_i^r(\omega; \beta_i^r, \gamma_{-i,\omega}^r)$ for $r = O, T$. Below, we examine reputational incentives under each regime and characterize best responses in detail. In our first result – which follows from Propositions 4.2 and 4.3 below – we focus on the common structure of best responses across regimes. We show that in both cases, best responses are in cutoff strategies; i.e., in each regime $r = O, T$ there exists a $\xi_i^r \in \mathbb{R}$ such that $\beta_i^r(s_i) = 0$ for all $s_i < \xi_i^r$ and $\beta_i^r(s_i) = 1$ for all $s_i > \xi_i^r$.

Proposition 4.1. *In each regime $r = O, T$, i 's best response is a cutoff strategy. As a result, $\sigma_i^r(\omega) = 1 - \Phi([\xi_i^r - \omega]\theta_i)$, where the cutoff ξ_i^r (uniquely) solves (4.4) with equality.*

Transparent Regime. In the Transparent regime, $\mathcal{Y}^T(v_i, v_{-i}) = \vec{v}$, and thus

$$\Psi_i^T(\omega; \beta_i^T, \gamma_{-i,\omega}^T) \equiv E_{v_{-i}} [P_i^T(0, v_{-i}; \beta_i^T, \gamma_{-i,\omega}^T) - P_i^T(1, v_{-i}; \beta_i^T, \gamma_{-i,\omega}^T)|\omega], \quad (4.5)$$

where $P_i^T(v_i, v_{-i}; \beta_i^T, \gamma_{-i,\omega}^T)$ denotes the Principal's belief that i is an expert after observing a recommendation profile (v_i, v_{-i}) given $(\beta_i^T, \gamma_{-i,\omega}^T)$. Since the Principal's inference about i will depend on the realization of other committee members' recommendations, i takes into consideration the entire distribution of v_{-i} , using $\Pr(\omega|s_i)$ and $\gamma_{-i,\omega}^T$. A key point to note here is that other members' strategies affect i 's problem exclusively by determining how much i 's Principal learns about the realization of the state ω from the observed recommendation profile v_{-i} , which in turn improves her ability to learn about i 's type. In this sense, other members' recommendations

are a signal about the realization of the state.²¹

Lemma D.1 in the Appendix characterizes $P_i^T(v_i, v_{-i}; \beta_i, \gamma_{-i}^T)$, alongside the Principal's posterior belief that $\omega = 1$ conditional on v_{-i} , $\tilde{\rho}_i(v_{-i}|\gamma_{-i}^T) \equiv \Pr(\omega = 1|v_{-i}, x_i; \gamma_{-i}^T)$. Intuitively, the Principal's posterior belief that member i is an expert when i recommends a dovish (hawkish) policy is *increasing* (*decreasing*) in $\sigma_i^T(1)$ and $\sigma_i^T(0)$, and $P_i(0, v_{-i}) \geq P_i(1, v_{-i})$ iff

$$\frac{\sigma_i(0)}{1 - \sigma_i(1)} \geq \frac{\tilde{\rho}(v_{-i}|\gamma_{-i}^T)}{1 - \tilde{\rho}(v_{-i}|\gamma_{-i}^T)}.$$

Since these results holds for all v_{-i} , it follows that $\Psi_i^T(\omega; \beta_i^T, \gamma_{-i}^T)$ is increasing in $\sigma_i^T(1)$ and $\sigma_i^T(0)$. In particular, when i adopts a cutoff strategy, i.e., $\sigma_i^T(\omega) = 1 - \Phi([\xi_i^T - \omega] \theta_i)$, the expected reputational gain from recommending a dovish vs a hawkish policy in state ω is decreasing in the cutoff ξ_i^T . In Lemma D.2, moreover, we show that $E_\omega[\Psi_i^T(\omega; \beta^T) | s_i]$ is decreasing in s_i . Since the left-hand side of (4.4) is increasing in s_i , it follows that player i 's best response must take the form of either a cutoff strategy or an unresponsive strategy. In the proof of Proposition 4.2, we rule out the latter and show that there exists a unique best response in cutoff strategies.

Proposition 4.2 (Best Responses in the Transparent Regime). *Agent i 's best response in the transparent regime is a cutoff strategy with cutoff ξ_i^T , uniquely given as the solution to*

$$\Pr(\omega = 1|\xi_i^T) \equiv \pi_i + \Delta_i E_\omega [\Psi_i^T(\omega; \xi_i^T, \gamma_{-i}^T) | s_i] |_{s_i = \xi_i^T} \quad (4.6)$$

Proof. See Appendix D □

Opaque Regime. In the Opaque regime, member i 's recommendation affects the Principal's observed outcome y only when i is pivotal; that is, when exactly $(n + 1)/2$ fellow members recommend a tight monetary policy. Let $t_i(\omega; \gamma_{-i}^O)$ denote the probability that i is pivotal in state ω , given others' conditional choice probabilities (CCPs) $\gamma_{-i}^O \equiv \{(\gamma_{j,0}^O, \gamma_{j,1}^O)\}_{j \neq i}$, and let $P_i^O(y; \beta_i^O, \gamma_{-i}^O)$ denote the Principal's posterior

²¹This differs from Ottaviani and Sørensen (2001) and Levy (2007). In these papers, there are no strategic interactions between agents in the Transparent regime, because the Principal is assumed to observe the realization of the state.

belief that i is an expert after observing $y \in \{0, 1\}$.²² Letting

$$\Psi_i^O(\omega; \beta_i^O, \gamma_{-i}^O) \equiv [P_i^O(0; \beta_i^O, \gamma_{-i}^O) - P_i^O(1; \beta_i^O, \gamma_{-i}^O)] t_i(\omega; \gamma_{-i}^O), \quad (4.7)$$

expression (4.4) can be written as $\beta_i^O(s_i) = 1$ if and only if

$$\Pr(\omega = 1 \mid s_i) \geq \pi_i + \Delta_i [P_i^O(0; \beta_i^O, \gamma_{-i}^O) - P_i^O(1; \beta_i^O, \gamma_{-i}^O)] \underbrace{\sum_{\omega} t_i(\omega; \gamma_{-i}^O) \Pr(\omega \mid s_i)}_{\Pr(\text{PIV}^i \mid s_i, \gamma_{-i}^O)}. \quad (4.3b)$$

Three observations are noteworthy at this point. First, differently to the Transparent regime, in the Opaque regime the principal learns about member i based on the committee's decision. For this, i has to be influential to the aggregate outcome. In fact, note that if $\Pr(\text{PIV}^i \mid s_i, \gamma_{-i}^O) \approx 0$, the reputational effect of the committee's decision on i becomes negligible. Thus, reputational effects matter for i only when other members are *divided* conditional on state ω .²³

Second, as it was the case in the Transparent regime, other members' strategies influence i 's best response solely via γ_{-i}^O , a result that is key to our estimation strategy. Furthermore, in Lemma D.3 we show that other members' strategies enter the posteriors $P_i^O(y; \beta_i^O, \gamma_{-i}^O)$ solely via (i) the pivotal probability $t_i(\omega; \gamma_{-i}^O)$, and (ii) the probability that the committee selects $y = 1$ in state ω independently of i 's action, $w_i(\omega; \gamma_{-i}^O)$. Thus, other members' strategies only affect i 's best response by how γ_{-i}^O affects $t_i(\omega; \gamma_{-i}^O)$ and $w_i(\omega; \gamma_{-i}^O)$.

Third, it follows immediately from Lemma D.3 that – as in the transparent regime – the Principal's posterior belief that member i is an expert when *the committee* sets on a dovish (hawkish) policy is *increasing (decreasing)* in $\sigma_i^O(1)$ and $\sigma_i^O(0)$. In particular, if i follows a cutoff strategy with threshold ξ_i^O , then $P_i^O(0) - P_i^O(1)$ is decreasing in ξ_i^O , and therefore $\Psi_i^O(\omega; \beta_i^O, \gamma_{-i}^O)$ is decreasing in ξ_i^O for $\omega = 0, 1$.

²²If, instead, formal voting dissents are observed by the public in the Opaque regime, the environment resembles the Transparent regime but with individual formal votes rather than policy recommendations, and with an added cost of dissent at the voting stage (see Appendix F). This cost captures the idea that once the policy directive is known from internal deliberation, members have incentives to conceal differences in their votes and project a united front to the public (Sibert (2003); Visser and Swank (2007)).

²³This rules out an equilibrium in the Opaque regime in which all members' decisions track the state closely (i.e., $\gamma_{j0}^O \approx 0$ and $\gamma_{j1}^O \approx 1$ for all $j \in \mathcal{N}$) and reputation has a large effect on behavior. In other words, if reputational considerations are to be meaningful in equilibrium, at least some members must make non-negligible mistakes.

Moreover, $P_i^O(0) \geq P_i^O(1)$ iff $\sum_{\omega} \rho(\omega) t_i(\omega; \gamma_{-i}^O) \sigma_i^O(\omega) \geq \rho t_i(1; \gamma_{-i}^O)$; i.e., if the probability that i 's recommendation changes the outcome from dovish to hawkish if she is *not* an expert is higher than the probability that it does when she *is* an expert. It follows that $P_i^O(0) > P_i^O(1)$ when i plays an unconditional hawkish strategy ($\sigma_i^O(\omega) = 1$ for all ω) and $P_i^O(0) < P_i^O(1)$ when i plays an unconditional dovish strategy ($\sigma_i^O(\omega) = 0$ for all ω).

We are now ready to state our next result. Note that (4.3b) can be written as

$$\beta_i^O(s_i) = 1 \Leftrightarrow \Pr(\omega = 1 \mid s_i) a(\sigma_i^O, \gamma_{-i}^O) \geq b(\sigma_i^O, \gamma_{-i}^O), \quad (4.3c)$$

where

$$a \equiv 1 - \Delta_i [P_i^O(0) - P_i^O(1)] (t_i(1) - t_i(0)), \quad b \equiv \pi_i + \Delta_i [P_i^O(0) - P_i^O(1)] t_i(0).$$

Since only $\Pr(\omega = 1 \mid s_i)$ depends on s_i , it follows immediately from (4.3c) that if β_i^O is a responsive strategy, it must be a cutoff strategy. Moreover, provided $a > 0$ (a mild assumption that holds at our estimates), any cutoff strategy must be increasing. Since our focus is on weakly increasing strategies, we henceforth assume that in equilibrium $|t_i(1; \gamma_{-i}^O) - t_i(0; \gamma_{-i}^O)| < 1/\Delta_i$ for all i . Finally, in the proof of our next result we show that any best response must be responsive, and hence a cutoff strategy.

Proposition 4.3 (Best Responses in the Opaque Regime). *Agent i 's best response in the Opaque regime is a cutoff strategy with cutoff ξ_i^O , uniquely given as the solution to*

$$\Pr(\omega = 1 \mid s_i) \big|_{s_i = \xi_i^O} \equiv \pi_i + \Delta_i E_{\omega} [\Psi_i^O(\omega; \xi_i^O; \gamma_{-i}^O) \mid s_i] \big|_{s_i = \xi_i^O} \quad (4.8)$$

Proof. See Appendix D □

Comparative Statics. Figure 3 plots the best-response cutoffs ξ_i^r in each regime, as functions of a member's bias π_i and reputational value Δ_i . The equilibrium cutoff ξ_i^r is increasing in π in both regimes. This is intuitive: ceteris paribus, a more hawkish agent cares more about the risk of inflation (adopting a dovish policy in a boom), and thus plays a more hawkish strategy, which favors the higher interest rate for a larger set of signals. On the other hand, we see in Figure 3 that ξ_i^r is decreasing in Δ in both regimes. This is not a general property. What is true is that, ceteris paribus, an agent with higher career concerns plays a more moderate strategy; i.e.,

$|\xi_i^r - 1/2|$ is decreasing in Δ . This is the *mimicking* incentive: an agent who puts a higher premium on reputation wants to look like an expert.




Figure 3: Best Response Cutoffs in the Opaque and Transparent Regimes. We compute best responses as π_i and Δ_i vary from their minimum to maximum estimated values. The simulation assumes a fixed committee of 12 members. We fix $\rho = 0.65$ and $\gamma_{1,j} = 0.8 = 1 - \gamma_{0,j}$ for all $j \neq i$.

As we can see in the left panel, the equilibrium cutoff ξ^r is more sensitive to π in the Opaque regime than in the Transparent regime. On the other hand, ξ^r is initially steeper in Δ under Transparency than in the Opaque regime, flattening out for larger values of Δ . To see the intuition for this result, let $F^r(\xi, \pi, \Delta) := S(\xi_i^r) - \Delta_i G(\xi_i^r) - \pi$, where $S(\xi_i^r) \equiv \Pr(\omega = 1 | s_i = \xi_i^r)$ and $G(\xi_i^r) \equiv E_\omega [\Psi_i^r(\omega | \xi_i^r; \gamma_{-i}^r) | s_i] |_{s_i = \xi_i^r}$. By the IFT, for each regime r ,

$$\frac{d\xi_i^r}{d\pi_i} = \frac{1}{S'(\xi_i^r) - \Delta_i G^{r'}(\xi_i^r)} \quad \text{and} \quad \frac{d\xi_i^r}{d\Delta_i} = \frac{G^r(\xi_i^r)}{S'(\xi_i^r) - \Delta_i G^{r'}(\xi_i^r)}$$

Intuitively, the principal's beliefs are more sensitive to the agent's action under transparency, where the action can be observed directly; i.e., $G^{T'}(\xi_i^T) < G^{O'}(\xi_i^O) < 0$. Since $S'(\xi_i^r) > 0$, and $\Delta > 0$, the denominator of both expressions is larger under Transparency. However, $d\xi_i^r/d\Delta_i$ depends on the level of the reputational effect, $G^r(\xi_i^r)$. And as we discussed before, the reputational gain from playing a dovish action is small when the agent has a dovish strategy (i.e., large ξ_i^T), and increases as ξ_i^T decreases

towards an unconditional hawkish strategy.

5 Identification and Estimation

5.1 Identification

Our identification strategy rests on two key ideas. First, we use the conditional choice probabilities of other agents under each regime, $\gamma_{-i,\omega}^r$, which can be computed directly from the data, to decompose the game into n independent decision problems. This decomposition is valid because, as shown in Section 4, other agents' best responses affect member i 's problem only through $\gamma_{-i,\omega}^r$. Second, we exploit the transparency shock, together with the availability of individual-level data in the Opaque regime, to compare equilibrium behavior under the same conditions across both regimes.²⁴

Consider a given committee membership, and (unlimited) data from given case characteristics X^a and X^b , observed in regimes $r = O, T$. Note that for $z = a, b$

$$\Pr(v|\psi^{r*}, r) = \prod_{t=1}^T \sum_{\omega} \rho^z(\omega) \prod_{i=1}^n (\gamma_{\omega i}^{rz})^{v_{it}} (1 - \gamma_{\omega i}^{rz})^{1-v_{it}}, \quad (5.1)$$

where $\gamma_{1i}^r \geq \gamma_{0i}^r$ for all $i \in N$ and $r = O, T$ since best responses are in increasing strategies in both regimes. The prior ρ^z , and the conditional choice probabilities $(\gamma_0^{rz}, \gamma_1^{rz})$ are then identified by standard arguments for mixture models (see Allman, Matias, and Rhodes (2009); Hall and Zhou (2003), Hu (2008), Kasahara and Shimotsu (2007)), provided $n \geq 3$. Denote the corresponding quantities obtained from the data as $\bar{\rho}^z$ and $\bar{\gamma}^{rz} \equiv \{(\bar{\gamma}_{j0}^{rz}, \bar{\gamma}_{j1}^{rz})\}_j$ for $r = O, T$. Since best responses are increasing cutoff strategies, for $z = a, b$ and $r = O, T$:

$$\begin{cases} \bar{\gamma}_{0,i}^{rz} = (1 - p_i) [1 - \Phi(\xi_i^{rz} \theta_i)] \\ \bar{\gamma}_{1,i}^{rz} = p_i + (1 - p_i) [1 - \Phi([\xi_i^{rz} - 1] \theta_i)] \end{cases} \quad (5.2)$$

²⁴Our identification strategy relies primarily on variation across cases, abstracting from changes in committee composition in the main analysis. Nonetheless, overlapping committee memberships provide additional information about the parameters of interest, which we incorporate into the estimation.

From these, we obtain

$$\frac{\bar{\gamma}_{0,i}^{ra}}{\bar{\gamma}_{0,i}^{wb}} = \frac{1 - \Phi(\xi_i^{ra}\theta_i)}{1 - \Phi(\xi_i^{wb}\theta_i)} \quad r, w \in \{O, T\} \quad (5.3)$$

and

$$\frac{\bar{\gamma}_{1,i}^{ra} - \bar{\gamma}_{1,i}^{wb}}{\bar{\gamma}_{0,i}^{ra} - \bar{\gamma}_{0,i}^{wb}} = \frac{\Phi([\xi_i^{wb} - 1]\theta_i) - \Phi([\xi_i^{ra} - 1]\theta_i)}{\Phi(\xi_i^{wb}\theta_i) - \Phi(\xi_i^{ra}\theta_i)} \quad r, w \in \{O, T\}, \quad (5.4)$$

from which we can recover the five unknowns $(\theta_i, \xi_i^{Ta}, \xi_i^{Tb}, \xi_i^{Oa}, \xi_i^{Ob})$. We can then use the first equation in (5.2) to get

$$p_i = \frac{1 - \Phi(\xi_i^{Ta}\theta_i) - \bar{\gamma}_{0,i}^{Ta}}{1 - \Phi(\xi_i^{Ta}\theta_i)}$$

Given $(\theta_i, p_i, \xi_i^{Oa}, \xi_i^{Ta})$, the pair (π_i, Δ_i) can be uniquely recovered from the equilibrium conditions (4.6) and (4.8) for member i in both regimes. Specifically, (π_i, Δ_i) solve the following system:

$$\begin{cases} \Pr(\omega = 1|\xi_i^O) = \pi_i + \Delta_i \sum_{\omega} \Psi_i^O(\omega; \xi_i^O, \gamma_{-i,\omega}^O) \Pr(\omega|\xi_i^O) \\ \Pr(\omega = 1|\xi_i^T) = \pi_i + \Delta_i \sum_{\omega} \Psi_i^T(\omega; \xi_i^T, \gamma_{-i,\omega}^T) \Pr(\omega|\xi_i^T) \end{cases}$$

giving us

$$\Delta_i = \frac{\Pr(\omega = 1|\xi_i^T) - \Pr(\omega = 1|\xi_i^O)}{\sum_{\omega} \{\Psi_i^T(\omega; \xi_i^T, \gamma_{-i,\omega}^T) \Pr(\omega|\xi_i^T) - \Psi_i^O(\omega; \xi_i^O, \gamma_{-i,\omega}^O) \Pr(\omega|\xi_i^O)\}} \quad (5.5)$$

and

$$\pi_i = \Pr(\omega = 1|\xi_i^O) - \Delta_i \sum_{\omega} \Psi_i^O(\omega; \xi_i^O, \gamma_{-i,\omega}^O) \Pr(\omega|\xi_i^O). \quad (5.6)$$

Remark 5.1 (Identification Intuition). *To develop an intuition about how the data maps to parameters, it is useful to assume p_i is known, and define*

$$\tilde{\gamma}_{0,i}^{rz} \equiv \frac{\bar{\gamma}_{0,i}^{rz}}{1 - p_i} = [1 - \Phi(\xi_i^{rz}\theta_i)] \quad \text{and} \quad \tilde{\gamma}_{1,i}^{rz} \equiv \frac{\bar{\gamma}_{1,i}^{rz} - p_i}{1 - p_i} = [1 - \Phi([\xi_i^{rz} - 1]\theta_i)].$$

Then we have

$$\theta_i = \Phi^{-1}(1 - \tilde{\gamma}_{0,i}^{rz}) - \Phi^{-1}(1 - \tilde{\gamma}_{1,i}^{rz}) \quad \text{and} \quad \xi_i^{rz} = \left[1 - \frac{\Phi^{-1}(1 - \tilde{\gamma}_{1,i}^{rz})}{\Phi^{-1}(1 - \tilde{\gamma}_{0,i}^{rz})} \right]^{-1},$$

as in Iaryczower and Shum (2012). Note that ability θ_i is increasing in the observed probability that i makes a correct recommendation in both booms $\Pr(v_i = 1|\omega = 1) = \tilde{\gamma}_{1,i}^{rz}$ and recessions, $\Pr(v_i = 0|\omega = 0) = 1 - \tilde{\gamma}_{0,i}^{rz}$. On the other hand, the equilibrium cutoff is pinned down by the asymmetry in the observed errors across states; $\xi_i^{rz} = 1/2$ if and only if i 's recommendations balance the errors in both states,

$$\Pr(v_i = 0|\omega = 1) \equiv 1 - \tilde{\gamma}_{1,i}^{rz} = \tilde{\gamma}_{0,i}^{rz} \equiv \Pr(v_i = 1|\omega = 0),$$

and ξ_i^{rz} increases (decreases), or equivalently is more dovish (hawkish) if and only if

$$\Pr(v_i = 0|\omega = 1) > (<) \Pr(v_i = 1|\omega = 0).$$

With this information, the equilibrium condition from one regime would allow us to identify a set of pairs (π_i, Δ_i) consistent with the data. The equilibrium conditions from both regimes allow us to disentangle bias π_i from career concerns Δ_i . $\square$

5.2 Estimation

Estimation follows the identification argument in two steps. In the first step, we obtain estimates for members' prior beliefs ρ and conditional choice probabilities for both the Opaque ($\gamma_{i\omega}^O$) and Transparent ($\gamma_{i\omega}^T$) regimes, as a function of macroeconomic fundamentals and member-specific covariates. In a second step, we recover structural parameters $\mu \equiv (\{\pi_i, \theta_i, \Delta_i\}_{i=1}^N, p, \tau)$ as a function of member-specific covariates. For simplicity, we show estimates with a common parameters p and τ across committee members.²⁵

First Stage Parameters. Given that equilibrium strategies are cutoff strategies, the first stage of our estimation procedure follows Iaryczower and Shum (2012), which we implement using the Bayesian approach of Iaryczower and Katz (2016). As in Iaryczower and Shum (2012), we treat the prior ρ_t in each meeting t and the conditional choice probabilities $\{\gamma_{\omega,it}^r\}$ as reduced-form objects, modeled as flexible functions of meeting-level covariates (X_t) and member-level covariates (Z_i) defined in Section 3. The vector of meeting characteristics X_t includes staff forecasts of in-

²⁵For robustness, we considered a fully heterogeneous model. We do not find p or τ to be significantly responsive to members' characteristics.

flation, real GDP growth, and unemployment, as well as recent money growth and the current level of the FFR. The vector of member characteristics Z_i includes indicators for gender and FOMC experience, compositional variables capturing career background, regional-bank affiliation dummies for Presidents, and indicators for the partisan affiliation of the appointing President for Governors.

Specifically, we assume that the prior ρ_t is a logistic function of macroeconomic fundamentals X_t observed by members in each meeting; i.e.,

$$\rho_t = \rho(X_t) = \frac{\exp(X_t' \beta)}{1 + \exp(X_t' \beta)}, \quad \beta \sim \mathcal{N}(0, \sigma_\beta), \quad \sigma_\beta \sim \mathcal{TN}(0, 2), \quad (5.7)$$

where $\mathcal{TN}(\cdot)$ denotes the truncated normal distribution. Conditional choice probabilities $\gamma_{0,it}$ and $\gamma_{1,it}$, in turn, are logistic functions of macroeconomic fundamentals X_t , member-specific covariates Z_i , and the regime type, $r \in \{O, T\}$. Letting $D_{it} = [X_t, Z_i]$,

$$\gamma_{i0}^r(D_{it}) = \frac{\exp(\alpha_{a[i]}^0 + D_{it}' \zeta^0)}{1 + \exp(\alpha_{a[i]}^0 + D_{it}' \zeta^0)} \quad \text{and} \quad \gamma_{i1}^r(D_{it}) = \frac{\gamma_{i0}(D_{it}) + \exp(\alpha_{a[i]}^1 + D_{it}' \zeta^1)}{1 + \exp(\alpha_{a[i]}^1 + D_{it}' \zeta^1)} \quad (5.8)$$

with

$$\alpha_{a[i]}^\omega \sim \mathcal{N}(0, \sigma_\alpha), \quad \zeta^\omega \sim \mathcal{N}(0, \sigma_\zeta), \quad \sigma_\alpha, \sigma_\zeta \sim \mathcal{N}(0, 2).$$

This parameterization in (5.8) imposes the constraint $\gamma_{1,it} \geq \gamma_{0,it}$, which ensures identification by resolving the label-switching problem. Letting $\iota \equiv (\beta, \zeta^0, \zeta^1, \alpha_{a[i]}^0, \alpha_{a[i]}^1)$, the likelihood is then

$$L(\mathbf{v}_t \mid \iota) \propto g(\iota) \times \prod_{t=1}^T \left\{ \rho(X_t) \prod_{i=1}^N \left[\gamma_{i1}(D_{it})^{v_{it}} (1 - \gamma_{i1}(D_{it}))^{1-v_{it}} \right] + (1 - \rho(X_t)) \prod_{i=1}^N \left[\gamma_{i0}(D_{it})^{v_{it}} (1 - \gamma_{i0}(D_{it}))^{1-v_{it}} \right] \right\}, \quad (5.9)$$

where $g(\iota)$ denotes the joint prior on the coefficients ι , assumed to be normally distributed with diagonal and heteroskedastic variance. We recover the posterior distribution of ι using an efficient Hamiltonian Monte Carlo algorithm, following Homan and Gelman (2014).²⁶ This procedure yields the posterior distribution of the prior

²⁶For estimation, we ran four parallel chains with dispersed initial values for 5,000 iterations each, discarding the first 3,000 as warm-up. Convergence for each parameter is assessed using the potential

belief and the conditional choice probabilities through equations (5.7) and (5.8).

Second Stage Parameters. In the second stage of the estimation, we combine the first-stage estimates with the model’s equilibrium conditions to recover the structural parameters μ , following the identification strategy outlined above. We allow the structural parameters $\{\pi_i, \theta_i, \Delta_i\}$ to depend on member characteristics Z_i as follows:

$$\pi_i = \frac{\exp(Z_i' \lambda_\pi)}{1 + \exp(Z_i' \lambda_\pi)}; \quad \theta_i = \exp(Z_i' \lambda_\theta); \quad \Delta_i = \exp(Z_i' \lambda_\Delta). \quad (5.10)$$

Given observable member characteristics Z , any parameter vector $\Lambda \equiv (\lambda_\pi, \lambda_\theta, \lambda_\Delta, p, \tau)$ implies a vector of structural parameters $\tilde{\mu}(\Lambda)$. Conditional on the estimated choice probabilities of other members, equilibrium conditions for each member $i \in N$ and meeting t yield theoretical conditional choice probabilities $\tilde{\gamma}_{0,it}^r$ and $\tilde{\gamma}_{1,it}^r$. These, in turn, define a pseudo-likelihood over the observed recommendation data.²⁷ Informally, our second-stage estimator selects the value of Λ that maximizes the pseudo-likelihood of members’ choices, treating $\{\gamma_{0,-i,t}^r\}_t$ and $\{\gamma_{1,-i,t}^r\}_t$ as given. To do so, we implement a nested algorithm with an outer loop that maximizes the likelihood of the data, and an inner loop that solves for the optimal cutpoints and generates best-response conditional choice probabilities for every member across meetings, as follows:

Inner Loop. For a given trial value $\tilde{\Lambda}$,

1. Use (5.10) to compute parameters $\{\pi_i(\Lambda), \theta_i(\Lambda), \Delta_i(\Lambda)\}$ for each $i \in N$.
2. For each member i and meeting t , use best responses (4.6) and (4.8) to compute equilibrium cutpoints $\xi_{it}^r(\Lambda)$ taking $\gamma_{0,-i,t}^r$ and $\gamma_{1,-i,t}^r$ as given.²⁸

scale reduction factor $\hat{R}$ Gelman and Rubin (1992) and by inspecting mixing across chains.

²⁷Aguirregabiria and Mira (2007) refer to the second-stage objective as a pseudo-likelihood because the choice probabilities used in this stage are not necessarily equilibrium probabilities given the structural parameters. Instead, they represent best responses to estimates of the population choice probabilities.

²⁸In practice, to make the computation feasible, we compute exact equilibrium cutpoints for every meeting in the sample (i.e., conditional on the prior $\tilde{p}_t$ and on other members’ behavior $\tilde{\gamma}_{-it,\omega}$) in a discrete grid of $S = 10,000$ points for the vector of structural parameters $\{\pi, \theta, \Delta, p, \tau\}$, and interpolate the equilibrium cutpoints for all other points in the parameter space via a flexible (meeting-specific) third-degree polynomial. We allow for different polynomial coefficients for Governors and Presidents to incorporate heterogeneous priors for the Presidents’ Principals, informed by regional macroeconomic shocks. Figure A5 in the Appendix shows the densities of cutpoints and the polynomial fit (correlation = 0.92).

3. Use $\xi_{it}^r(\Lambda)$ and $\theta_i(\Lambda)$ to compute $\sigma_i^r(\omega; \Lambda) = [1 - \Phi([\xi_{it}^r(\Lambda) - \omega]\theta_i(\Lambda))]$, and then

$$\gamma_{i,\omega}^r(\Lambda) = p(\Lambda)\omega + (1 - p(\Lambda))\sigma_i^r(\omega; \Lambda)$$

Outer Loop. Our estimate of $\vec{\Lambda}$ is then

$$\vec{\Lambda}^* \equiv \arg \max_{\vec{\Lambda}} \prod_{t=1}^T \sum_{\omega} \rho_t(\omega) \prod_{i=1}^N [\gamma_{i,\omega}^{r(t)}(\Lambda)]^{v_{it}} [1 - \gamma_{i,\omega}^{r(t)}(\Lambda)]^{1-v_{it}} \quad (5.11)$$

For inference, we use the asymptotic covariance matrix recovered from the Hessian inverse at the maximum likelihood estimates.

Goodness of Fit. Appendix E reports several metrics evaluating the model’s fit to the data. The model correctly predicts about 70% of individual policy recommendations. A likelihood ratio test shows that incorporating career concerns significantly improves the fit relative to a purely expressive model ($\Delta = 0$). Moreover, the structural model successfully replicates the observed patterns of disagreement among FOMC members documented in the reduced-form evidence, both in policy recommendations and in the content of members’ speeches, even though these variables were not directly targeted in estimation.

6 Results

In this section we present our results. We begin with the estimates of the structural parameters $\mu \equiv (\{\pi_i, \theta_i, \Delta_i\}_{i=1}^N, p, \tau)$, including the prior function $\rho(\cdot)$. We then turn to the main results on the effects of career concerns and transparency, in Section 6.1.

Structural Parameter Estimates. The left panel of Figure 4 plots the estimated coefficients of the prior function, showing the effect of each covariate on the prior belief that the economy is in a boom. All coefficient estimates have the expected sign: higher expected GDP growth, higher expected inflation, lower expected unemployment, and a lower prevailing federal funds rate are all associated with a higher probability that the economy is in a boom. By contrast, conditional on expected inflation, the money supply (M1) has no statistically significant effect on prior beliefs. The right panel of Figure 4 depicts the evolution of the common prior over time. The recovered

prior closely tracks economic conditions, exhibiting sharp declines during recessions (indicated by grey bars). The figure also highlights that the 1970s and 1980s were marked by greater volatility, with large fluctuations in prior beliefs that are not observed in later periods.




Figure 4: Prior Probability Estimates. Left: Coefficient estimates for the ρ function. The figure plots the effect of changing each covariate from its minimum to its maximum observed value in the sample. Thin (thick) lines representing the 95% (90%) credible intervals. Right: Estimate of Prior Belief $\rho = \Pr(\omega = 1)$ across meetings. Solid lines represent the posterior median, shaded region represents the 95% credible interval.

Figure 5 presents the estimates for reputation (Δ), individual competence (θ), and bias (π). The top panel reports the median value of each parameter estimate across the sample, separately for Governors and Presidents. The bottom panel plots the empirical distribution of estimates across members in the sample.²⁹

The left column of Figure 5 plots our estimates of members' career concerns; that is, the value Δ a member assigns to her reputation relative to the policy decision under consideration. As the top panel shows, the typical Governor places greater weight on reputation than the typical President (although confidence intervals overlap, a difference-in-means test confirms this result.) In fact, a meeting-level analysis reveals that the meeting-specific median Governor consistently places more weight on reputation than the meeting-specific median President. As the bottom panel shows,

²⁹Figure A3 in the Appendix presents the marginal effect of each individual-specific covariate on the estimates of (Δ, θ, π) , as well as affiliation fixed effects.




Figure 5: Estimates of Structural Parameters. Top panel: estimates of the median value of reputation (Δ), precision (θ), and bias (π) across members and meetings, for Governors and Presidents. Thick (thin) lines represent 90% (95%) confidence intervals. Bottom panel: empirical distribution of the parameter estimates across members in the sample. Thin (thick) lines plot the interquartile range (90th percentile) of the empirical distribution.

however, there is substantial overlap in the empirical distributions of Δ for Governors and Presidents. Notably, rookie members assign a higher value to establishing a reputation for competence (see Figure A3). This finding is consistent with stylized facts in the literature, and supports the analysis of Hansen, McMahon, and Prat (2018), who proxy career concerns with members' experience. Similarly, members with government experience or central bank experience tend to assign a higher value to their reputation, while Female members generally put a higher value on getting the decision right.

The middle and right panels plot our estimates of members' information precision (θ) and bias (π). The results generally confirm previous findings in the literature. We

find that Governors are relatively dovish ($\pi > 1/2$), while Presidents are relatively hawkish ($\pi < 1/2$), and that the typical Governor has more precise information than the typical President (López-Moctezuma (2016)). As it was true for Δ , though, there is substantial heterogeneity within each group: 16% of Fed Presidents are more dovish than the median Governor. Similarly 16% of Presidents have better information than the median Governor, reflecting the characteristics of newly appointed members over time. In particular, both Governors and Presidents were relatively hawkish in the mid-1990s, but became more dovish around the time of the Great Recession in 2008. All else equal, rookie members tend to have less precise information than their more senior counterparts, and tend to be slightly more dovish than their senior counterparts (see Figure A3).

Figure A4 in the Appendix reports estimates of the prior probability that a member is of the expert type, p , and of the precision of regional unemployment as a signal of the state in principals' beliefs, τ . The prior on members' type is tightly estimated around 0.5, while the precision of regional shocks is precisely estimated around 0.3.

6.1 Main Results

Having presented our estimates, we jump to our main results. We begin by describing the effectiveness of committee recommendations given the economic conditions and committee compositions observed in the data. We then isolate the effect of reputational considerations, and evaluate counterfactual policy experiments.

Outcomes in the Data. To evaluate the effectiveness of decision-making in the FOMC, we focus on the probability that FOMC members give a correct recommendation. In particular, we focus on two main outcomes: (i) the conditional probability of a correct recommendation in booms and recessions:

$$\gamma_{1,i,t} = \Pr(v_{it} = 1 | \omega_t = 1) \quad \text{and} \quad 1 - \gamma_{0,i,t} = \Pr(v_{it} = 0 | \omega_t = 0),$$

and (ii) the ex-ante probability of a correct recommendation:

$$\bar{\gamma}_{it} \equiv \rho_t \gamma_{1,i,t} + (1 - \rho_t)(1 - \gamma_{0,i,t}).$$

The ex-ante probability of a correct recommendation provides a benchmark mea-

sure of effectiveness, as it weights outcomes across states by the probability with which each state is realized. By contrast, the state-contingent probabilities are crucial for understanding the payoff consequences of equilibrium outcomes, since different Principals can, and typically do, weight errors differently across states. A hawkish Principal ($\pi^P < 1/2$) is primarily concerned with inflation and therefore more willing to tolerate mistakes of failing to lower rates during recessions in order to reduce the incidence of excessively loose monetary policy during booms. Conversely, a dovish Principal ($\pi^P > 1/2$) places greater emphasis on mitigating recessionary risks and is thus more willing to accept mistakes of not raising rates during booms in order to minimize the frequency of contractionary policy errors in recessions.




Figure 6: Conditional Choice Probabilities for Presidents and Governors. Probability of a Correct Recommendation in booms (γ_1), recessions ($1 - \gamma_0$) and ex ante ($\bar{\gamma}$) for the median President and Governor in each meeting. Solid lines represent the median estimate, shaded region represents the 95% confidence interval.

Figure 6 plots the conditional probability of a correct recommendation in booms and recessions, as well as the ex-ante probability, for the median Governor and President in each meeting. The dashed vertical line marks the regime change, while gray bars denote recessions. Both Governors and Presidents have relatively low error rates in booms (left column), especially under Transparency. In particular, Presidents – who as we saw are typically more hawkish than Governors – tend to make fewer mistakes in booms across both regimes. By contrast, errors in recessions are more pronounced. This is particularly so for Presidents, who tend to favor overly cautious (i.e., contractionary) policies when downturns are realized. The right panel of Fig-

ure 6 shows the ex-ante probability of a correct recommendation $\bar{\gamma}$ for the median Governor and President, weighting state-contingent probabilities by the prior ρ_t in each meeting. When the likelihood of each state is accounted for, Governors consistently achieve higher ex-ante accuracy than Presidents, in both regimes and across economic conditions.

Career Concerns and Counterfactual Regimes. We now turn to our main questions: how career concerns shape behavior within a given institutional regime, and how regimes affect behavior through their impact on career incentives. To address these, we begin by comparing outcomes *with* career concerns (as in the data) with the outcomes that would emerge *without* career concerns (i.e., setting $\Delta = 0$ for all members). We then evaluate a counterfactual in which the regime is held fixed throughout the sample.




Figure 7: Conditional Choice Probabilities for Presidents and Governors. Probability of a Correct Recommendation in booms (γ_1), recessions ($1 - \gamma_0$) and ex ante ($\bar{\gamma}$) for the median President and Governor in each meeting. Solid lines represent the median estimate, shaded region represents the 95% confidence interval.

Figure 7 presents the results of the first exercise. In both regimes, and for the vast majority of meetings, career concerns increase the median error in booms and reduce the median error in recessions. In the Opaque regime, however, the effect on the probability of a correct recommendation is relatively small, while in Transparency the impact is substantial. We should be clear that these gains and losses are *potential* gains and losses, which are only materialize if the state of the economy actually

realizes as a recession or a boom. Still, they are relevant to inform the preferred institutional configuration of a principal who prioritizes the risk of recessions versus inflation or vice versa. The third column shows that when we weight the expected gains and losses by ρ , career concerns have a moderately positive effect on the *ex ante* probability of a correct recommendation.

Figure 8 presents outcomes under a counterfactual scenario in which the institutional regime is held constant throughout the sample. Specifically, we plot the probability of a correct recommendation in booms and recessions under a Transparent regime for the entire sample (i.e., counterfactual under the Opaque regime) and under an Opaque regime for the entire sample (i.e., counterfactual under Transparency). As shown in the figure, Transparency typically outperforms Opacity in recessions, while the Opaque regime dominates in booms. However, when considering the *ex ante* probability of a correct recommendation, Transparency consistently dominates across the entire sample. In other words, for an unbiased Principal – who values policy errors in recessions and booms equally – Transparency yields to better outcomes throughout the sample period.




Figure 8: Data and Counterfactual Conditional Choice Probabilities. Probability of a correct recommendation in booms, recessions and *ex ante* under a Transparent regime in the entire sample (counterfactual in the Opaque regime) and under an Opaque regime in the entire sample (counterfactual in the Transparent regime).

Mechanisms: Pandering and Moderation. In this section, we analyze the mechanisms driving the difference in individual behavior with and without career

concerns, captured in our framework by the change in equilibrium cutpoints:

$$\text{Total} \equiv \xi_{\pi}^{noCC} - \xi_{\pi}^{CC}. \quad (6.1)$$

A naive explanation might attribute this difference entirely to pandering (or anti-pandering), whereby a committee member biases her actions toward (or against) the prior relative to the decision she would have taken based solely on her information, absent career concerns. This view, however, is incomplete. The notion of pandering / anti-pandering applies cleanly only when the agent is unbiased ($\pi = 1/2$), as is commonly assumed in the theoretical literature on reputation for competence. When $\pi \neq 1/2$, the difference conflates pandering with the incentive to mimic the expert type. To disentangle these two forces, we define pandering relative to the behavior of an agent's *unbiased analog*. In particular, we ask whether an unbiased version of the agent tilts her behavior toward (or against) the prior relative to the decision she would have taken absent career concerns:

$$\text{Pandering} = (\xi_{1/2}^{noCC} - \xi_{1/2}^{CC}) \mathbb{I}(\rho - \tfrac{1}{2}).$$

We define mimicking as the component of behavior driven by career concerns that reflects the difference between a biased agent ($\pi \neq 1/2$) and her unbiased counterpart. Formally, we compute it as the residual between:

1. the change induced by both preferences and career concerns, $A \equiv \xi_{1/2}^{CC} - \xi_{\pi}^{CC}$, and
2. the change induced by preferences alone, $B \equiv \xi_{1/2}^{noCC} - \xi_{\pi}^{noCC}$.

Hence,

$$\underbrace{\xi_{\pi}^{noCC} - \xi_{\pi}^{CC}}_{\text{Total}} = \underbrace{(\xi_{1/2}^{noCC} - \xi_{1/2}^{CC})}_{\text{Pandering}^*} + \underbrace{(\xi_{1/2}^{CC} - \xi_{\pi}^{CC})}_{\text{Mimicking}} - \underbrace{(\xi_{1/2}^{noCC} - \xi_{\pi}^{noCC})}_{\text{Mimicking}}, \quad (6.2)$$

where Pandering* isolates the change in cutoffs potentially attributable to pandering/anti-pandering.

Figure 9 presents this decomposition. For each term in equation (6.2), the figure plots the distribution of the corresponding outcomes for the median President and




Figure 9: Behavioral Decomposition. We decompose the total change in behavior between career concerns and no career concerns, $\xi_{\pi}^{\text{noCC}} - \xi_{\pi}^{\text{CC}}$, into a *pandering* and a *mimicking* effect. The figures plot the distribution of outcomes for the median President and median Governor in each meeting.

median Governor in each meeting, under the Opaque (left) and Transparent (right) regimes. The top panel shows the total change in behavior due to career concerns, $\xi_{\pi}^{\text{noCC}} - \xi_{\pi}^{\text{CC}}$. The second row reports the (anti-)pandering effect, $\xi_{1/2}^{\text{noCC}} - \xi_{1/2}^{\text{CC}}$. The third row shows the extent of mimicking, defined as the residual between two components: (i) the change in behavior induced by both preferences and career concerns for a biased agent relative to an unbiased one, $A \equiv \xi_{1/2}^{\text{CC}} - \xi_{\pi}^{\text{CC}}$ (plotted in the fourth row), and (ii) the change induced by preferences alone, $B \equiv \xi_{1/2}^{\text{noCC}} - \xi_{\pi}^{\text{noCC}}$ (plotted in the fifth row).

In the Opaque regime, career concerns have only a modest, generally hawkish, effect on equilibrium strategies, driven almost entirely by anti-pandering. As the last row (term B) shows, in the absence of career concerns, preferences alone would generate a substantial difference in behavior relative to that of an unbiased agent.

Yet the change in behavior *with* career concerns (term A) is virtually identical to that *without* them. In short, under Opaqueness, the difference in behavior between a biased and an unbiased agent stems almost entirely from preferences rather than from reputational incentives.

In the Transparent regime, by contrast, career concerns have a large and heterogeneous effect on equilibrium strategies, driven by both anti-pandering and mimicking. As in the Opaque regime, we observe a pro-hawkish anti-pandering effect, though it is stronger under Transparency, particularly for Presidents. The more salient difference, however, stems from mimicking. As before, preferences alone would generate a sizable difference in behavior (term B). But under Transparency, this difference is substantially muted: members strategically moderate their behavior to mimic the expert type, with the extent of adjustment determined by their career concern parameter Δ_i . As a result, hawkish members adopt more dovish strategies than they would absent career concerns, while dovish members adopt more hawkish strategies.

These results have important implications for understanding the role of career concerns within the FOMC. A cursory reading of the theoretical and empirical literature might suggest that Transparency worsens policy outcomes by inducing anti-pandering and conformity. Our findings challenge this view. While career concerns do significantly affect equilibrium outcomes, the dominant mechanism is the moderation of biases through mimicking expert behavior. This strategic moderation offsets heterogeneity in policy preferences within the committee and – provided the principal is sufficiently moderate – can improve welfare.

Conformity Revisited: Model-Based Heterogeneous Treatment Effects.

Next, we use our estimates to quantify the effect of transparency on conformity – long the primary focus of the theoretical and empirical literature studying career concerns in committees.

We first define (gross) conformity between member i and the Chair as the probability that they make the same recommendation:

$$C_i = \sum_{\omega} \rho(\omega) \left[\gamma_{i\omega} \gamma_{ch,\omega} + (1 - \gamma_{i\omega}) (1 - \gamma_{ch,\omega}) \right]. \quad (6.3)$$

The measure in (6.3) corresponds to the conventional usage of the term in the literature, which conflates strategic behavior with the correlation in members' actions

arising from interdependent values. To isolate the *strategic* component of conformity, we compute a counterfactual measure of conformity under the assumption that agents have no career concerns, denoted C_i^{noCC} :

$$C_i^{\text{noCC}} = \sum_{\omega} \rho(\omega) \left[\gamma_{i\omega}^{\text{noCC}} \gamma_{ch,\omega}^{\text{noCC}} + (1 - \gamma_{i\omega}^{\text{noCC}}) (1 - \gamma_{ch,\omega}^{\text{noCC}}) \right],$$

where $\gamma_{i\omega}^{\text{noCC}}$ are the counterfactual conditional choice probabilities computed under the restriction $\Delta_i = 0$ for all i . We then define *strategic conformity*, $\bar{C}_i$, as the difference between gross conformity and this baseline level due purely to interdependent values:

$$\bar{C}_i \equiv C_i - C_i^{\text{noCC}}. \quad (6.4)$$

The left panel of Figure 10 plots gross conformity (C_i) for the median Governor and President across regimes. The point estimates are consistent with the before and after specification (model (1) in Table 1), showing an increase of conformity in Transparency for both Governors and Presidents, but do not align with the DiD result that transparency increases Governors' conformity relative to Presidents (the effects are imprecisely estimated).




Figure 10: Conformity and Strategic Conformity. Measures of Conformity with respect to the chair, for the median Governor and President across regimes. Thick (thin) lines represent 90%(95%) confidence intervals.

Differently to the case of gross conformity, the figure for strategic conformity ($\bar{C}_i$) in the right panel shows a clear distinction across regimes, for both Governors and Presidents. The model-based estimate of the DiD effect is consistent with the reduced-form results, showing that transparency increases Governors' conformity relative to Presidents. But differently to the reduced-form estimates from the “before and after” specification – which can fail to appropriately control for changing economic conditions (Hansen, McMahon, and Prat (2018)) – the model-based estimate shows that Transparency *reduces* conformity for Presidents.

To further understand the connection between the model-based and reduced-form estimates of the DiD effect, we use the model to compute heterogenous treatment effects. Note that the estimate of the β coefficient in regression (3.1) can be written in terms of potential outcomes as:

$$\hat{\beta} = (E[y_{it} | P, T] - E[y_{it} | G, T]) - (E[y_{it} | P, O] - E[y_{it} | G, O]),$$

where $E[y_{it} | a, r]$ denotes expected disagreement conditional on appointment type $a \in \{P, G\}$ and regime $r \in \{T, O\}$. Within the model, expected disagreement can be written as

$$E[y_{it} | a, r] = \sum_{\omega} \rho_t(\omega) \left[\gamma_{\omega it}^{ar} (1 - \gamma_{\omega ct}^r) + (1 - \gamma_{\omega it}^{ar}) \gamma_{\omega ct}^r \right] = 1 - C_{it}^{ar},$$

where C_{it}^{ar} is the (gross) conformity measure for appointment type a under regime r . This expression can be computed counterfactually for any meeting t by solving for potential outcomes given the structural parameter estimates. Using this, we construct a model-based analog of the DiD effect at the meeting level, using the difference in potential outcomes:

$$\tilde{\beta}_{it} = (C_{it}^{GT} - C_{it}^{GO}) - (C_{it}^{PT} - C_{it}^{PO}) \tag{6.5}$$

$$= (\bar{C}_{it}^{GT} - \bar{C}_{it}^{GO}) - (\bar{C}_{it}^{PT} - \bar{C}_{it}^{PO}), \tag{6.6}$$

where $\bar{C}_{it}^{ar}$ denotes strategic conformity for appointment type a and regime r .

Figure 11 shows the empirical distribution of the median value of $\tilde{\beta}_{it}$ for each meeting t , denoted $\tilde{\beta}_{m(t)}$. The green dashed line indicates the point estimate from regression (3.1), while the purple dashed line marks the average of $\tilde{\beta}_{m(t)}$ across meet-




Figure 11: Model-based DiD Effect. Distribution of the median value of $\tilde{\beta}_{it}$ for each meeting t , $\tilde{\beta}_{m(t)}$. The green dashed line shows the point estimate from equation (3.1). The yellow dashed line shows the average value of $\tilde{\beta}_{m(t)}$ across meetings.

ings. The average transparency effect from the structural estimates (0.043) is close to the reduced-form DiD estimate (0.059). More importantly, the model-based DiD approach recovers potential outcomes at the individual level and quantifies heterogeneity by affiliation – an analysis not feasible with reduced-form methods alone. On average, transparency increases the probability of conformity with the Chair for Governors relative to Presidents by less than 0.06. Yet the effect is more than double this magnitude in more than 10% of meetings, while transparency induces greater conformity among Presidents than Governors in 10% of the meetings.

6.2 Counterfactual: Heterogeneity of Regional Shocks

In the reduced-form analysis we showed that Fed Presidents are more responsive to regional shocks under Transparency than in the Opaque regime. This suggests Fed Presidents are at least partially attuned to local principals. To capture this channel, in the model we allowed principals to have heterogeneous priors ρ_i , which are informed by local economic conditions (the difference between regional and national unemployment). Our estimate of the variance of the informational content of regional

shocks in Principals’ beliefs, ζ , supports this channel. Recall that

$$\rho_i \equiv \Pr(\omega = 1|x_i) = \left[1 + \left(\frac{1 - \rho}{\rho} \right) \exp(2\zeta^2 x_i) \right]^{-1},$$

so that the homogeneous priors case requires $\zeta \rightarrow 0$. Instead, our estimate of $\hat{\zeta} \approx 0.3$ – with a 95% confidence interval above 0.25 – indicates that Fed Presidents at least partially target regional Principals’ interests. This in turn indicates that both conformity and the probability of a correct recommendation depends on the distribution of regional shocks.

To explore this channel further, we use our estimates to assess the impact of regional shocks on *equilibrium* outcomes. Specifically, we compare observed behavior to two counterfactuals, for which we recompute equilibrium behavior at each meeting. In the first, we zero out all regional shocks. In the second, we amplify them by 50%. To reduce computational demands, we simulate nine-member committees composed of six Presidents and three Governors; the Chair and two randomly drawn Governors from each meeting.




Figure 12: Effect of Regional Shocks. Conformity (left) and probability of a correct recommendation (right), in the data and in each counterfactual treatment. For each treatment (“observed”, “no shocks”, and “amplified”), the figure plots the distribution of the median value of the outcome for each meeting.

We present the counterfactual results in Figure 12. The left panel plots conformity, while the right panel the probability of a correct recommendation, for the observed data and each counterfactual treatment. For each case (“observed”, “no shocks”, and “amplified”), the figure shows the distribution of the median outcome across meetings. The results indicate that greater heterogeneity in regional shocks has no effect under Opacity, but reduces conformity *and* increases the likelihood of policy “errors” (with respect to the agents’ prior ρ) in the Transparent regime.

To assess more directly the impact of regional shocks on Fed Presidents’ behavior and the effectiveness of policy recommendations, we simulate outcomes under two counterfactual committees defined by affiliation. In one, all members are Fed Presidents, whose behavior is influenced by local economic conditions; in the other, all members are Board Governors, who are not disproportionately affected by any particular region. This design isolates the role of affiliation under alternative committee compositions. For both counterfactuals, we fix aggregate conditions by setting the common prior ρ to its sample median and hold committee size constant at seven members (the actual number of Governors). Committees are simulated by randomly sampling individuals, conditional on affiliation, from the full set of members observed in the data.

Figure 13 presents the results. Committees composed entirely of Presidents are less cohesive than those composed of Governors, a gap that widens under Transparency, where the probability of conformity falls by 4 p.p. relative to Opacity. Committee composition also produces significant differences in policy effectiveness: all-Governor committees consistently generate more accurate recommendations than all-President committees, and this accuracy gap expands under Transparency. Because Presidents are more exposed to regional shocks in the Transparent regime, the quality of their recommendations declines relative to Governors.

7 Conclusions

In this paper, we use a structural estimation approach and individual data on FOMC members’ policy recommendations to quantify the distortions induced by career concerns in a collective decision-making body. The restrictions implied by the model, together with the unanticipated regime shift from Opacity to Transparency, uniquely identify the model parameters and allows us to separately recover individual-level pa-




Figure 13: Conformity (left) and probability of a correct recommendation (right) under two counterfactual committee compositions: (i) all members are Fed Presidents (“President”), and (ii) all members are Board Governors (“Governor”). We set ρ to its sample median, and fix committee size at seven members – the actual number of Governors – drawing randomly from the pool of FOMC participants conditional on affiliation.

rameters that capture members’ preferences and reputation for competence.

A cursory reading of the theoretical and empirical literature on career concerns might suggest that Transparency worsens policy outcomes by inducing anti-pandering and conformity. Our results challenge this view. We find that career concerns can improve the accuracy of policy recommendations. In the Transparent regime, where the effects of career concerns on behavior are stronger, Transparency delivers superior outcomes. While career concerns do generate anti-pandering in Transparency, its dominant effect is to moderate biases through mimicking expert behavior. This mechanism offsets heterogeneity in preferences and can enhance welfare, even when conformity rises for some members.

References

- ADOLPH, C. (2013): *Bankers, bureaucrats, and central bank politics: The myth of neutrality*. Cambridge University Press.
- AGUIRREGABIRIA, V., AND P. MIRA (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1–53.
- ALLMAN, E. S., C. MATIAS, AND J. A. RHODES (2009): “Identifiability of parameters in latent structure models with many observed variables,” *Annals of Statistics*, 37(6A), 3099–3132.
- ARUOBA, S. B., A. DRAZEN, AND R. VLAICU (2018): “A structural model of electoral accountability,” *International Economic Review*, <https://doi.org/10.1111/iere.12361>.
- BAJARI, P., L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BANKS, J., AND R. SUNDARAM (1998): “Optimal Retention in Agency Problems,” *Journal of Economic Theory*, 82(2), 293–323.
- BLINDER, A. S. (2008): *The quiet revolution: Central banking goes modern*. Yale University Press.
- CANEN, N. J., AND M. IARYCZOWER (2024): “Innovation Adoption by Committee: Evaluating Decision-Making in the FDA,” Discussion paper, National Bureau of Economic Research.
- CANES-WRONE, B., M. HERRON, AND K. SHOTTS (2001): “Leadership and Pandering: A Theory of Executive Policymaking,” *American Journal of Political Science*, 45(3), 532–550.
- CHAPPELL, H., R. MCGREGOR, AND T. VERMILYEA (2005): *Committee decisions on monetary policy: Evidence from historical records of the Federal Open Market Committee*. The MIT Press.
- CHAPPELL JR, H. W., T. M. HAVRILESKY, AND R. R. MCGREGOR (1997): “Monetary policy preferences of individual FOMC members: a content analysis of the memoranda of discussion,” *Review of Economics and Statistics*, 79(3), 454–460.
- CHAPPELL JR, H. W., R. R. MCGREGOR, AND T. A. VERMILYEA (2008): “Regional economic conditions and monetary policy,” *European Journal of Political Economy*, 24(2), 283–293.
- DIERMEIER, D., M. KEANE, AND A. MERLO (2005): “A Political Economy Model of Congressional Careers,” *American Economic Review*, pp. 347–373.
- DRISCOLL, J. C., AND A. C. KRAAY (1998): “Consistent covariance matrix estimation with spatially dependent panel data,” *Review of economics and statistics*, 80(4), 549–560.

- EFFINGER, M. R., AND M. K. POLBORN (2001): “Herding and anti-herding: A model of reputational differentiation,” *European Economic Review*, 45(3), 385–403.
- FINAN, F., AND M. MAZZOCCO (2021): “Electoral incentives and the allocation of public funds,” *Journal of the European Economic Association*, 19(5), 2467–2512.
- GELMAN, A., AND D. B. RUBIN (1992): “Inference from iterative simulation using multiple sequences,” *Statistical science*, pp. 457–472.
- GERLACH-KRISTEN, P. (2004): “Is the MPC’s voting record informative about future UK monetary policy?,” *Scandinavian Journal of Economics*, 106(2), 299–313.
- GILDEA, J. A. (1992): “The regional representation of Federal Reserve Bank presidents,” *Journal of Money, Credit and Banking*, 24(2), 215–225.
- GOODFRIEND, M. (1999): “The role of a regional bank in a system of central banks,” in *Carnegie-Rochester Conference Series on Public Policy*, vol. 51, pp. 51–71. Elsevier.
- HALL, P., AND X.-H. ZHOU (2003): “Nonparametric Estimation of Component Distributions in a Multivariate Mixture,” *Annals of Statistics*, 31, 201–224.
- HANSEN, S., AND M. MCMAHON (2016): “First impressions matter: Signalling as a source of policy dynamics,” *Review of Economic Studies*, 83(4), 1645–1672.
- HANSEN, S., M. MCMAHON, AND A. PRAT (2018): “Transparency and deliberation within the FOMC: a computational linguistics approach,” *The Quarterly Journal of Economics*, 133(2), 801–870.
- HANSEN, S., M. MCMAHON, AND C. V. RIVERA (2014): “Preferences or private assessments on a monetary policy committee?,” *Journal of monetary Economics*, 67, 16–32.
- HARRIS, M. N., AND C. SPENCER (2009): “The policy choices and reaction functions of Bank of England MPC members,” *Southern Economic Journal*, 76(2), 482–499.
- HAYO, B., AND M. NEUENKIRCH (2013): “Do Federal Reserve presidents communicate with a regional bias?,” *Journal of Macroeconomics*, 35, 62–72.
- HIX, S., B. HØYLAND, AND N. VIVYAN (2010): “From doves to hawks: A spatial analysis of voting in the Monetary Policy Committee of the Bank of England,” *European Journal of Political Research*, 49(6), 731–758.
- HOLMSTRÖM, B. (1982): “Managerial incentive problems: A dynamic perspective,” in *Essays in Economics and Management in Honor of Lars Wahlbeck*. Swedish School of Economics, Helsinki.
- HOMAN, M. D., AND A. GELMAN (2014): “The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo,” *The Journal of Machine Learning Research*, 15(1), 1593–1623.
- HOTZ, V. J., AND R. A. MILLER (1993): “Conditional choice probabilities and the estimation of dynamic models,” *The Review of Economic Studies*, 60(3), 497–529.

- HU, Y. (2008): “Identification and Estimation of Nonlinear Models with Misclassification Error Using Instrumental Variables: a General Solution,” *Journal of Econometrics*, 144, 27–61.
- IARYCZOWER, M., AND G. KATZ (2016): “More than politics: Ability and ideology in the British Appellate Committee,” *The Journal of Law, Economics, and Organization*, 32(1), 61–93.
- IARYCZOWER, M., G. LEWIS, AND M. SHUM (2013): “To Elect or to Appoint? Bias, Information, and Responsiveness of Bureaucrats and Politicians,” *Journal of Public Economics*, 97, 230–244.
- IARYCZOWER, M., G. LOPEZ-MOCTEZUMA, AND A. MEIROWITZ (2024): “Career Concerns and the Dynamics of Electoral Accountability,” *American Journal of Political Science*, 68(2), 696–713, Typset, Princeton University.
- IARYCZOWER, M., X. SHI, AND M. SHUM (2018): “Can words get in the way? The effect of deliberation in collective decision making,” *Journal of Political Economy*, 126(2), 688–734.
- IARYCZOWER, M., AND M. SHUM (2012): “The Value of Information in the Court. Get it Right, Keep it Tight,” *American Economic Review*, 102, 202–237.
- KASAHARA, H., AND K. SHIMOTSU (2007): “Nonparametric Identification and Estimation of Multivariate Mixtures,” Queen’s University, Department of Economics.
- LEVY, G. (2004): “Anti-herding and strategic consultation,” *European Economic Review*, 48(3), 503–525.
- (2007): “Decision making in committees: Transparency, reputation, and voting rules,” *American economic review*, 97(1), 150–168.
- LIM, C. S. (2013): “Preferences and incentives of appointed and elected public officials: Evidence from state trial court judges,” *The American Economic Review*, 103(4), 1360–1397.
- LINDSEY, D. E. (2003): “A modern history of FOMC communication: 1975-2002,” *The Federal Reserve Board of Governors*.
- LÓPEZ-MOCTEZUMA, G. (2016): “Sequential deliberation in collective decision-making: The case of the FOMC,” *Unpublished manuscript*.
- MASKIN, E., AND J. TIROLE (2004): “The politician and the judge: Accountability in government,” *American Economic Review*, 94(4), 1034–1054.
- MEADE, E. E. (2005): “The FOMC: preferences, voting, and consensus,” *Federal Reserve Bank of St. Louis Review*, 87(March/April 2005).
- MEADE, E. E., AND D. N. SHEETS (2005): “Regional influences on FOMC voting patterns,” *Journal of Money, Credit and Banking*, pp. 661–677.

- MEADE, E. E., AND D. STASAVAGE (2008): “Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve,” *The Economic Journal*, 118(528), 695–717.
- MORRIS, S. (2001): “Political Correctness,” *Journal of Political Economy*, 109(2), 231–265.
- OTTAVIANI, M., AND P. SØRENSEN (2001): “Information aggregation in debate: who should speak first?,” *Journal of Public Economics*, 81(3), 393–421.
- OTTAVIANI, M., AND P. N. SØRENSEN (2006): “Professional advice,” *Journal of Economic Theory*, 126(1), 120–142.
- PRAT, A. (2005): “The Wrong Kind of Transparency,” *American Economic Review*, 95(3), 862–877.
- PRENDERGAST, C., AND L. STOLE (1996): “Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning,” *Journal of Political Economy*, 104(6), 1105–1134.
- ROBERTS, M. E., B. M. STEWART, D. TINGLEY, E. M. AIROLDI, ET AL. (2013): “The structural topic model and applied social science,” in *Advances in neural information processing systems workshop on topic models: computation, application, and evaluation*, vol. 4, pp. 1–20. Harrahs and Harveys, Lake Tahoe.
- SCHARFSTEIN, D. S., AND J. C. STEIN (1990): “Herd Behavior and Investment,” *The American Economic Review*, 80(3), 465–479.
- SIBERT, A. (2003): “Monetary policy committees: individual and collective reputations,” *The Review of Economic Studies*, 70(3), 649–665.
- (2009): “Is transparency about central bank plans desirable?,” *Journal of the European Economic Association*, 7(4), 831–857.
- SIEG, H., AND C. YOON (2017): “Estimating dynamic games of electoral competition to evaluate term limits in us gubernatorial elections,” *American Economic Review*, 107(7), 1824–57.
- STASAVAGE, D. (2007): “Polarization and publicity: rethinking the benefits of deliberative democracy,” *The Journal of Politics*, 69(1), 59–72.
- SWANK, J., B. VISSER, AND O. H. SWANK (2008): “How committees of experts interact with the outside world: some theory, and evidence from the FOMC,” *Journal of the European Economic Association*, 6(2-3), 478–486.
- VISSER, B., AND O. SWANK (2007): “On Committees of Experts,” *Quarterly Journal of Economics*, 122, 337–372.
- ZWIEBEL, J. (1995): “Corporate conservatism and relative compensation,” *Journal of Political economy*, 103(1), 1–25.